

PERCEPTION

Howard Robinson



THE PROBLEMS OF PHILOSOPHY
edited by Ted Honderich

**Also available as a printed book
see title verso for ISBN details**

Perception

Perception

Howard Robinson



London and New York

First published 1994
by Routledge
11 New Fetter Lane, London EC4P 4EE

This edition published in the Taylor & Francis e-Library, 2003.

Simultaneously published in the USA and Canada
by Routledge
29 West 35th Street, New York, NY 10001

© 1994 Howard Robinson

All rights reserved. No part of this book may be reprinted or reproduced or utilized in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloging in Publication Data
Robinson, Howard.

Perception/Howard Robinson.

p. cm.—(Problems of philosophy)

Includes bibliographical references and index.

1. Perception (Philosophy) I. Title. II. Series: Problems of
philosophy (Routledge (Firm))

B828.45.R62 1994

121'.3—dc20 93—49381

ISBN 0-203-16435-0 Master e-book ISBN

ISBN 0-203-25848-7 (Adobe eReader Format)
ISBN 0-415-03364-0 (Print Edition)

In memory of my parents,
Benjamin and Ethel Robinson

Contents

| | |
|---|-----------|
| Preface | xi |
| I The Classical Empiricist Conception of the Content of Perceptual Experience | 1 |
| 1 The traditional empiricist conception of sensory content | 1 |
| 2 The ancient and medieval background to the empiricist conception of perception | 4 |
| 3 Descartes and the empiricist conception | 11 |
| 4 Locke and the traditional conception | 11 |
| 5 Reasons for the disappearance of the intentional interpretation | 14 |
| 6 Intentionality from Reid to Husserl | 19 |
| 7 The percept theory as an alternative to the classical conception | 27 |
| 8 Conclusion | 30 |
| II The Traditional Arguments for the Empiricist Conception of Sense-contents: the Argument from Illusion | 31 |
| 1 The basic argument and the Phenomenal Principle | 31 |
| 2 Alternative accounts of appearance | 44 |
| 3 The difficulty of finding an alternative to the Phenomenal Principle | 54 |
| 4 Generalising the argument | 56 |
| 5 Conclusion | 57 |
| III Further Arguments against Naive Realism | 59 |
| 1 The argument from secondary qualities | 59 |
| 2 Argument from science stated | 74 |

Contents

| | | |
|-----------|---|-----|
| 3 | The paradigm case argument | 76 |
| 4 | Simple compatibilism | 77 |
| 5 | Perspectivalism | 77 |
| 6 | Dispositional direct realism and the manifest image of the world | 79 |
| 7 | The time-lag argument | 80 |
| 8 | The causal argument | 84 |
| 9 | The argument from hallucination | 87 |
| 10 | Conclusion of Chapters I–III | 89 |
| IV | Sense-data and the Anti-private Language Argument | 91 |
| 1 | Varieties of the argument | 91 |
| 2 | What the argument is meant to prove | 92 |
| 3 | The traditional interpretation of the anti-private language argument | 95 |
| 4 | The verificationist attack on the recognitional conception | 99 |
| 5 | Kripke's radical interpretation: the outline | 105 |
| 6 | The problem with the infinite content of a rule: Wittgenstein's treatment | 107 |
| 7 | Infinite content: Kripke's statement of the problem | 109 |
| 8 | Kripke's treatment of the dispositional solution | 110 |
| 9 | Kripke and the intellectualist solution | 111 |
| 10 | The social solution | 113 |
| 11 | Walker's refutation of the social solution | 115 |
| 12 | Not an argument but an alternative account of language | 116 |
| 13 | Memory, 'seeming' and privacy: an appendix to the traditional argument | 117 |
| V | Contemporary Physicalist Theories of Perception | 119 |
| 1 | Experience and reductionism | 119 |
| 2 | Pure cognitivism | 122 |
| 3 | The perceptual core | 125 |
| 4 | Physicalism and externalism | 136 |
| 5 | Externalism, knowledge and perception | 139 |
| 6 | Martin's objections | 140 |
| 7 | The 'straw man' objection | 141 |
| 8 | Macdonald's weak externalism | 145 |
| 9 | External to the mind or the head? | 147 |

| | | |
|-------------|--|-----|
| VI | The Revised—and Successful—Causal Argument for Sense-data | 151 |
| 1 | The argument stated | 151 |
| 2 | The attack on (2): the disjunctive analysis | 152 |
| 3 | Anti-realist intuitions | 158 |
| 4 | Defence of premise (1) | 159 |
| 5 | Conclusion | 162 |
| VII | The Intentional and Adverbial Theories | 163 |
| 1 | Motivation | 163 |
| 2 | Intentional objects and psychological states | 164 |
| 3 | Grammar and intentionality | 167 |
| 4 | The ‘as if’ interpretation of intentionality | 172 |
| 5 | Adverbialism and naïve realism | 174 |
| 6 | Adverbial content and sensory object | 177 |
| 7 | The refutation of adverbialism | 180 |
| 8 | The history of adverbialism | 182 |
| 9 | <i>Esse est percipi</i> and being an object of consciousness | 184 |
| 10 | Conclusion | 185 |
| VIII | The Nature of Sense-data | 187 |
| 1 | Possible positions on the nature of the common element | 187 |
| 2 | Arguments for phenomenal indeterminacy | 190 |
| 3 | Dennett and the central case: the knowledge argument | 198 |
| 4 | The situation so far | 202 |
| 5 | The maximal position and the blind spot | 202 |
| 6 | The maximal position and depth | 205 |
| 7 | Spatial transmodality: Molyneux’s problem | 207 |
| 8 | Conclusion | 212 |
| IX | Sense-data and the Physical World | 213 |
| 1 | The options | 213 |
| 2 | Varieties of representative realism | 213 |
| 3 | General objections to representative realism | 215 |
| 4 | Types of phenomenalism | 226 |
| 5 | The Berkeleian possibility | 229 |
| 6 | Humean, non-analytic reductionism | 231 |
| 7 | Phenomenalism and common sense | 236 |
| 8 | Conclusion | 238 |

Contents

| | |
|--------------|-----|
| Notes | 239 |
| Bibliography | 251 |
| Index | 258 |

Preface

This book is a history and a defence of what, with some shorthand, I will call the sense-datum theory of perception. It is part of the rubric for books in the series *The Problems of Philosophy* that a section of the book should be historical and another section should tell the plain truth, unburdened by scholarship. I have followed this requirement, at least roughly, giving the history of the empiricist theory of sensory contents and the traditional arguments for it in Chapters I to III, and devising and defending somewhat more satisfactory arguments and investigating the nature of sense-data and their relation to the external world in Chapters VI to IX. Chapters IV and V, however, can either be regarded as the contemporary end of the history—the supposed final nails in the coffin of the sense-datum theory—or as current errors that must be dispatched before the truth can be properly defended. The upshot is meant to be the vindication of the sense-datum theory in a fairly traditional form, and, more tentatively, a preference for the phenomenalist rather than representationalist account of the data's relation to the physical world. This latter eccentricity is, however, quite inessential to the main conclusion in defence of sense-data.

My interest in the philosophy of perception began in a small discussion group at school, an offshoot of the main philosophy society (happily named the Berkeley Society), which was, like the Berkeley Society, guided and inspired by John Armstrong. Since then I owe a great deal to many with whom I have argued on this subject, including Christopher Taylor, Jonathan Barnes, Lesley Brown, Penelope Mackie, Paul Snowdon, Michael Martin; to many discussions presided over by Ralph Walker, and to years of argument with John Foster. Some of the work in the book was read to the interdisciplinary 'Liverpool Philosophy of Mind Group'. Part or all of the manuscript was read and

Preface

usefully criticised by Siobh  n Phillips, Nicholas Nathan and Barry Dainton. The work was greatly assisted by two periods of leave granted me by my university.

CHAPTER I

The Classical Empiricist Conception of the Content of Perceptual Experience

1 The traditional empiricist conception of sensory content

The majority of modern philosophers—that is, the majority of philosophers writing since the seventeenth century—have believed that in perception one is aware of some item other than the physical object one takes oneself to be perceiving. So if I see a tree, for example, what I am *really* or *directly* aware of is something which can roughly be thought of as a tree-image, in or before my mind, rather than a mind-independent physical tree in external physical space. The *ideas* of Locke and Berkeley, Hume's *impressions* and the *qualia*, *sensa* and *sense-data* of twentieth-century philosophers are all generally supposed to be of this type. On the other hand, the majority of strictly contemporary philosophers—that is, the majority of philosophers active in the analytic tradition since the Second World War—have denied that one need postulate such entities and affirm that we are, normally, directly aware of the external world itself.

The earlier modern conception I shall call the *classical* or *empiricist conception of sense-contents*, and this book is largely about the arguments that are brought for and against this conception. Contrary to the general contemporary opinion, I shall be arguing that this classical empiricist conception is essentially correct—that there are things more or less of the kind which I shall, for convenience, call 'sense-data'. The empiricist conception of sense-contents—sense-data—can be roughly characterised as follows.¹

A sense-datum, as I shall understand it, meets five conditions:

- 1 It is something of which we are aware.
- 2 It is non-physical.
- 3 Its occurrence is logically private to a single subject.

- 4 It actually possesses standard sensible qualities, for example, shape, colour, loudness, 'feel' of various sorts.
- 5 It possesses no intrinsic intentionality; that is, though it may suggest to the mind through habit other things 'beyond' it, in itself it possesses only sensible qualities which do not refer beyond themselves.

The concept of intentionality plays such a central role in anti-sense-datum theories that a brief introduction to it at the outset is essential. Intentionality is the property of states—usually mental states—of being *about* things. An ordinary physical state is not about anything, it is just *there*, unless it is *interpreted* as being about or signifying something. But thoughts, beliefs, desires and most mental attitudes are essentially about things—they refer beyond themselves; they are directed upon objects. These objects need not be actual things, for one can think about, believe in, desire, etc., things that do not exist. The objects of intentional verbs are called 'intentional objects', and, because they need not exist for the activity to be directed to them, they are described as 'intentionally *inexistent*'. If perceptual states are intentional then they are essentially *of* things, as are thoughts. In this case the question will not arise of how perceptual states can put us in touch with the external world, for they are essentially of things in the external world. Furthermore, as the objects of intentional activities can be things that do not exist, we would not need to worry about the status of things that we seem to perceive that are not really there. On the other hand, intentional states are mysterious things. If they are intrinsically about other things, what properties, if any, do they possess intrinsically? Sense—data are, in some ways, less mysterious than intentional states, for the properties sense-data are supposed to possess are straightforward sensible qualities. These issues will recur persistently in this book.

The outright opponent of sense-data denies that there are any private objects of awareness involved in normal perception. A more moderate opponent allows that there are private objects of awareness, but claims that their content is purely intentional; that is, that all the content is essentially—if only seemingly—of external physical features and objects, the contents themselves not actually possessing or instantiating any sensible properties. A still more moderate opponent would allow that there are private objects of awareness that instantiate certain qualities—at least including these secondary qualities—but also

possessing intentional properties, so that a patch of brown could be essentially as of a table-shaped public object, quite independently of any inclination of the subject so to interpret it. In my conclusions I shall tend to favour the full sense-datum theory, but will remain somewhat agnostic about the presence of a moderate intentional element.

There are two different views of the relationship between sense-data and the physical world. According to idealists and phenomenologists, such as Berkeley, Hume and Mill, the data are actually part of physical objects, for objects consist only of actual or actual and possible sense-data; according to representative realists, such as Locke is generally supposed to be, they are caused by physical objects, which they resemble, though the resemblance may only be a rather abstract and structural one. I shall discuss representationalism and phenomenism in Chapter IX.

Many further questions arise concerning sense-data even for their proponents; for example, whether they are mental or neither mental nor physical; whether in the visual case they are two or three dimensional; whether knowledge of them is incorrigible and, in general, what their relation is to concept-involving activities such as recognition; whether they are particulars or a sort of sensory universal; whether it is coherent, within the scope of the classical conception, to regard them as modes of sensory activity rather than objects of it. Some of these issues will be dealt with in the following pages of this book. In Chapters II and III, my main purpose is to consider the arguments that have traditionally been brought for the classical conception minimally conceived in accordance with (1)–(5) above, and the objections raised against these arguments. The arguments take the form of attempted refutations of naive or direct realist theories of perception, the thought being that if one is not directly aware of external objects, then one must instead be aware of sense-data as defined by our five conditions. There is obviously some slack here, because one could reject naive realism without adopting exactly the classical conception in one's replacement theory: the sense-datum theory is merely the commonest candidate. One could, for example, hold, with Russell, that we are aware of states of our own brain; or, with certain phenomenologists, that the private data do possess intrinsic intentionality.² These and other variations will come up in the course of the discussion, which is centred on the classical conception.

In the remainder of this chapter I shall look, not at the arguments for sense-data, but, briefly, at the role of the empiricist conception in the history of philosophy. I do not try to be definitive in this, but to raise some of the questions, largely about intentionality and perception, that will recur throughout the book.

2 The ancient and medieval background to the empiricist conception of perception

It is natural to make a connection between the empiricist conception and earlier philosophies via the notion of *phantasm*. In the Aristotelian tradition, phantasms are images set up in the senses and, among the scholastics at least, they look very much like a kind of sense-datum.³ The interest of classical theories in a study like this, however, is not to provide primitive cases of empiricism, but to bring out certain tensions that exist in the project of developing an adequate theory of perception. In particular they can be used to illustrate the problems that arise when one tries to harmonise cognition with sensation. I shall not, therefore, begin the discussion of classical theories with phantasms, but start with what looks like a bizarre attempt to reconcile the human *activity* of perception with the *passive* reception of sensations.

Certain early Greek philosophers claimed that vision occurs when a stream of particles coming from an object meets a physical emission coming from the eye.⁴ As what we know of these theories is rather limited, we cannot be sure why they opted for such a seemingly odd theory, but it can be interpreted in a way that helps to set up one of the basic problems of perception. This problem can be expressed as follows. On the one hand, a little reflection—that is, thought that does not resort to any science that goes beyond common experience—shows that perception involves some sort of physical influence running from the external object to the sense organ of the perceiver. On the other hand, the essential nature of experience seems to be that the subject mentally reaches out to, and makes conscious contact with, the external object. The directionality of the physical process and that of the lived experience seem to be in direct conflict. How can a process in which the subject is the passive recipient of a stimulus be the physical aspect or realisation of a process in which the subject reaches actively and consciously out into the world? The way that perception both reaches out from subjects and reaches into

them expresses itself in three tensions or dichotomies. These are between the physical process and the conscious or phenomenological aspects of perception; between the fact that experience and its immediate causes are internal—‘in the head’—on the one hand, and that its content is apparently external, on the other; and the tension between the passive receptivity and interpretative activity of the subject in perception. Perhaps it is helpful to think of the bizarre theory of effluents as an unconscious attempt to resolve these tensions. The active mental contribution is represented as a *physical* outflow, and the experience is somehow caused ‘out there’—if not where the object perceived is located, at least, as a compromise, half-way there and not in the eye or the brain. Early Greek philosophers are often said to be scientists rather than philosophers, and this dual causal process can be seen as an attempt to integrate in a purely physical story the active role of consciousness with the physical action of the stimulus. This way of reconciling the three tensions is, of course, blatantly inadequate. Empirically it fails because there is no evidence of a flow out from the senses; conceptually it fails because there is nothing to explain why such a collision of particles should constitute an experience. It seems to present a solution to the problem only by failing to grasp the more subtle way in which the causal and the phenomenological levels are both interdependent yet different in kind.

The usual way of reconciling these dilemmas has been to say that the physical process causes a representation or picture of the world *within* the subject and that the subject *interprets this as being a feature of*, or *projects it onto*, the external world. This is, roughly, the representative theory of perception which is found in Locke and is adopted by most moderate empiricists. As a way of solving the problem it has two features which might be thought of as shortcomings. First, it seems to be an irredeemably dualistic account, for, although the process whereby the object influences the subject is, in principle, an unproblematic physical process, the projecting and interpreting of images are thoroughly mentalistic: so there is an uncomfortable change of gear in the middle of the process of perception. Second, and more important, the projection or interpretation of an internal image is not *really* a reaching out to the external world, only the illusion of doing so.

The most important ancient and medieval discussions of perception take place under the influence of Aristotle, and Aristotle’s theory can

be seen as an attempt to reconcile the apparently conflicting factors without moving to a simple representative theory. Considered as a physical process, perception involves the perturbation of the sense organ so as to produce an image or phantasm which is—or is more or less—physical. Considered philosophically—that is, in the light of the need to give an account of what it is to perceive, as opposed to an account of how perception works—perception is *the reception of the form of an object without its matter*. So if I see Socrates, his form is present in a special way in the sense. The image set up in the organ is the vehicle for the form of the object perceived, so that through the physical process the very thing perceived (and not just a replica of it) is present to the perceiving subject.⁵ In this way the inward flow of the causal process and the outward reaching of experience are reconciled, and we see how the latter supervenes on the former.

Two basic questions arise about this theory. The first question is historical and scholarly: is it a correct account of Aristotle? A distinction could be made here between Aristotle and his followers, especially St Thomas Aquinas and the scholastic tradition. Some interpreters (as we shall see) think that Aristotle himself had a straightforwardly materialistic account of perception, but that Aquinas had developed it in the direction of the theory just sketched. The second question is whether such a theory can succeed in reconciling the causal facts and the common-sense phenomenology.

The simplest and most reductive interpretation of Aristotle says that (1) the image involved in perception is a purely physical object, so that when someone sees yellow the eye literally becomes yellow and that is the image that constitutes the sensation; (2) the *form without the matter* is the same thing as the image, because ‘without matter’ means no more than that the external object itself (or bits of it) does not enter the organ. So perception is a purely physical process in which a straightforwardly physical quality is established in the sense organ. This is transmitted from thence into the heart, where the senses are united, and is retained in the bloodstream, thereby constituting memory. Whether or not this is what Aristotle meant, he was usually interpreted by his commentators and successors in a less materialistic way.⁶ Most importantly, ‘form without matter’ was given a more mysterious interpretation, so that it signified a special kind of presence in the subject of the thing perceived. The scholastics expressed this by saying that the form of the object perceived possessed *esse intentionale*, not *esse materiale*—intentional existence, not material

existence—in the sense. Brentano interpreted this as meaning that the form is an intentional object and exists in the soul, as mental, and not in the body, as physical.⁷ It is the possibility of this interpretation that makes the Aristotelian theory relevant to a general treatment of perception, for the notion of an intentional object will keep recurring throughout this book. It is a mysterious notion in certain fundamental respects, and the Aristotelian-cum-scholastic theories provide material for an interesting interpretation of it. The intentional state is actually constituted by a mode of existence of its object; so there is no mystery about how the mental content relates to the external reality, for it *is* that reality—or the form of that reality—taking on a psychological mode of existence.

Unfortunately, Brentano's theory is not merely disputed as an interpretation of Aristotle, for there are also those who think it too metaphysical even as a reading of Aquinas. Such sceptics think that the notion of intentional existence in Aquinas has a less interesting sense than this interpretation gives to it. Sheldon Cohen has argued that, according to Aquinas, there is nothing particularly psychological about intentional existence.⁸ A colour exists intentionally in a mirror or in the air, because in neither case is the thing coloured in the ordinary, material sense. So if one looks at a red object in a mirror, the mirror is not literally red, in the same sense as that in which the object reflected is red, although red is, in some sense, in the mirror. And when the form of red passes from an object to the eye, the air in between does not become red. Intentional existence is not necessarily, therefore, mental. The jelly of the eye becomes red in just the sense in which a mirror does; the image in the eye has the same peculiar ontological status as a reflection, if one tries to take reflections realistically.

The example of the mirror does not, however, naturalise or physicalise the notion of intentional existence in the way Cohen thinks. The case of the mirror is not incompatible with the principal thrust of Brentano's theory. In what sense is the form of red in the mirror, if not literally? Only, I think, in the sense that there is something in the mirror which communicates, transmits or makes available, the red in the object. This is exactly what the intentional object does in the psychological version of the theory. The only difference is that, in the case of the mirror, there need not be anything to transmit it to, for no one might be looking, whereas, in the case of a sense, the subject is automatically present. It is true that

Brentano thinks of the intentional object as mental, whereas, on this theory, intentionality can be embodied; but its logical features remain the same, namely those of communicating, transmitting or making available its object, without possessing the quality of that object in a literal material manner, as would a simple pictorial representation. It would be fair to say that the whole problem for 'intentional object' theories of perception is whether sense can be made of the idea that there is an entity that can do this; that is, represent its content *transparently*, not constituting the kind of 'veil of perception' that the empiricists' contents are supposed to constitute. Whether such 'transparent' vehicles of representation are immaterial or physical is an extra—though important—point.

A concept is, of course, something that is essentially *of* its object, and so represents it transparently; and Aristotle talks of the intellect—which is wholly immaterial—as receiving form without matter, for there is no matter in the intellect.⁹ This might seem to make thought and perception very similar. But even if one is prepared to accept concepts as things that represent their content transparently, this does not solve the problem for perception. Some philosophers wish to treat perception simply as the acquisition of beliefs, so that perceptual content is not essentially different from purely conceptual content. But, in fact, something extra is needed to give the presentational nature of perception. This is something that will be discussed later, but Aristotle clearly believes that there is a crucial difference: that is why perception involves images in a physical sense organ, whilst intellect cannot be embodied. The image in the mirror may be just like the form in the eye or the faculty of sight, but it is not just like the form in the intellect. And the difference explains the experiential quality of perception. The analogy with a mirror image may seem to solve the problem of reconciling intentionality and sensory content, because such a thing possesses both the sensory element—by being an image—and the intentionality, because the image is essentially of its object and does not actually possess the property of the object, as would a literal picture. For this to work one would have to accept the scholastic account of mirror images, as the intentional yet physical manifestation of the form of the object. There has to be, that is, some special vehicle involved, corresponding to the intentional form in the mirror. If one adopts a simple causal story of how perception works, with light rays from the object to the eye being reflected in the mirror and causing an experience in the subject, then no particular

illumination is given about the nature of experience by the example of the mirror; for, according to the causal account, red is not present in any sense—intentional or otherwise—in the mirror: the mirror simply reflects the light. If the Aristotelian way of blending the sensible and the intentional rests on finding extra-mental cases in nature which somehow make clear what it is for a form to exist intentionally—without matter—then the attempt seems to be a failure, for these cases are either not really cases of matterless forms or are just as mysterious as the psychological case.

There is another way in which one might hope that an Aristotelian theory might overcome the problems of representative realism. What comes into the sense embodied in the image is the *form of the very object perceived*. It might, therefore, seem that, even if the image is no better than a sense-datum, the fact that it carries the formal identity of the object with it makes that object present in a stronger sense than is possible on a normal representative realism. If it really is Socrates' humanity in my sense then I really am in direct contact with him. Whether this comes to anything significant depends on what sense one can make of the idea that it is the form of the object perceived that enters the sense.

The expression 'the form of the object perceived' can be interpreted in either of two ways. One might observe that forms are *kinds* of things and conclude that having the form of an object in the sense means the same as having a form of *the same kind* as the object's in the sense. So having the form of Socrates in the sense will just mean having *a* form of humanity in the sense. It is difficult to see how, if the thing in the sense is merely *of the same kind* as the object out there, it can constitute an especially good way of being in touch with *that very object* as opposed to having a matching kind of experience. So on this interpretation of 'form of the object' the problems of representationalism do not seem to be alleviated.

On the other hand, one might try to capitalise on the doctrine of particularised forms, and say that the form in the sense is not just a form of the same kind as that in the object, but the same particular form with a different mode of existence. The doctrine of particular forms states that forms are not universals but particulars. Thus Socrates' humanity is a different individual from Aristotle's humanity, rather than being the same universal existing in both. Whether this is Aristotle's account of form is controversial, but it is certainly the account accepted by St Thomas and the scholastic

tradition.¹⁰ On this account it might seem that the same individual form that was in the object perceived might be in the sense, and not merely some replica: it is the thing itself, in a sensory intentional mode. There is certainly pressure towards this sort of realism about mental content in the Aristotelian system. Aristotle identifies thought and its objects and this merely carries on that principle. To say that thought is identical with its object is equivalent to saying that the content of an act of thinking is the thing thought of, in a psychological mode. If good sense could be made of this idea then the gulf between representation and its object could be overcome. Unfortunately, in the case of pure thought its objects are forms and general truths, not particular objects, and the idea that a concept is the psychological manifestation of a form or general proposition is easier to swallow than any such identification of a sensible form and its object—unless, that is, one were to be a naive realist. The naive realist does identify perceptual content with perceived features of the external world and, to that extent, makes experience—or, at least, its contents—identical to its objects. A case can be made for interpreting Aristotle as a naive realist. On this view, the idea that form is received without matter means only that the object is taken in cognitively and sensibly without its having to enter the subject physically. The form is not, in some more or less literal sense, *transmitted* to the perceiver, passing through the medium between them; the form without the matter is just the object *qua* perceived. I shall put aside this interpretation, not because a textual case cannot be made for it, but because it deprives Aristotle of his use to us. Our interest in his theories rests on the hope that he can reconcile causal input and conscious outreach, but naive realism places its emphasis entirely on the latter; it represents one of the horns of the dilemma, not a reconciliation. The story which has an intentional entity—the form—abstracted from a physically transmitted image was at least a candidate for bringing about this reconciliation. The task facing intentional theories in general is to explain how a content caused in the subject can constitute a not-merely-pictorial grasp on the object. The hope that the supposedly Aristotelian way of explaining intentionality might do this looks as if it is vain, for we have not found an interpretation of ‘form without matter’ in perception that can give transparency to sensory representations.

3 Descartes and the empiricist conception

The standard view—in that sense of ‘standard’ which implies a good deal of over-simplification and even caricature—was that the empiricist conception of content is contemporaneous with modern philosophy: that the ideas of Descartes—at least as they relate to perception—are sense-data. The development from scholastic theory was taken to be as follows. ‘Modern philosophy’ consisted in a rejection of Aristotelian talk of forms and species, in perception as in other matters: so what survived from past theories was the phantasm, but conceived as mental rather than physical: a Cartesian or Lockean idea is, in perception at least, just a ‘spiritualised’ phantasm.

No serious scholar ever thought that Descartes’ position was as straightforward as this. He makes it quite clear that ideas possess what he calls ‘objective reality’, which means that it is part of their essential nature to have an object—that is, to be *of* something.¹¹ This contradicts the fifth condition of the empiricist conception, as described above, which was that sense-contents possess no intrinsic intentionality. Descartes is operating within a scholastic framework, in so far as he is retaining these intentional conceptions, but is trying not to depend essentially on the traditional notions of form or species. Descartes makes a move towards the modern position by affirming that intentionality must operate through some intermediary entity or representation—an idea—and is not just the object, or the form of the object, present to the mind. The difference here may seem insubstantial: is not the form in the mind a sort of intermediary entity? Yes and no. It is not a different thing from the formal nature of the object, though that nature is taking on a different sort of existence. Talk of ‘form’ suggests a pure intentionality—that is, it suggests that the representation is wholly transparent, possessing no other property than that of being *of* its object. Talk of ‘ideas in the mind’ lead one towards thinking about the intrinsic nature of the representation, and endowing it with properties in its own right. As we shall see, in the case of perception at least, these begin to get in the way of its purely representational role.

4 Locke and the traditional conception

Descartes’ explicit use of the scholastic concept of ‘objective reality’ to signify intentionality makes it clear that he is not adopting a simple

empiricist conception of content. Locke's studied refusal to say anything much that is illuminating about his use of 'idea' beyond that ideas are 'whatsoever is the object of a man's understanding whenever a man thinks', makes it almost impossible to attempt a fruitful interpretation of his intentions.¹² This is no doubt because Locke is specifically determined not to become involved in scholastic problems, as is evidenced by the continuation of the above quotation: 'I have used it to express whatever is meant by *phantasm*, *notion*, *species* or whatever it is which the mind can be employed about in thinking'.¹³ Locke does not employ the concepts of phantasm or species (in this sense) in his system, nor does he have any corresponding concepts. 'Idea' is, therefore, not simply generic for a range of crucially different sorts of mental content; rather Locke shows no sign of thinking that interesting and difficult theoretical differences are to be found within the category of ideas. This does not, of course, mean that he does not attach significance to different types of content, as represented by simple and complex ideas, modes, etc., but that *qua* ideas they are of the same kind. In particular, the crucial difference between sensory content—phantasm—and intellectual content—species—is not treated as if it marked a type-difference of a sort which raises important questions. In a way which would be difficult to credit if scholarship had not so singularly failed to show the opposite in Locke, and if the texts of Berkeley and Hume did not so plainly confirm it, the empiricists appear to be totally oblivious of the existence of any serious problem concerning the intentionality of thought and the relation between the contents and objects of mental states. The flight from the problems of the Schools has about it a wilful blindness. None of Locke, Berkeley or Hume shows any signs of serious thinking about the relation of their concepts of quality, idea or impression to the problem of universals: it is as if they thought that this issue had disappeared once one had eliminated talk of 'substantial forms' from science. Similarly, Berkeley in his nominalist theory of thought invokes the notion of one idea standing for a group without showing any sign that the notion of *standing for* is seriously in need of illumination.¹⁴ Finally, Hume is reduced to saying that something is an idea of—a concept of—something else just by being a faded replica of it.¹⁵ Locke does not say such things but his determined refusal to recognise that the scholastic theories of thought were concerned with a real issue played a vital role in making respectable the frame of mind which led to such doctrines.

In Locke's case, if there is a rationalisation of his failure to face problems of intentionality and generality, other than hostility to scholasticism, it might consist in a failure to distinguish epistemological from analytical questions. Suppose we take 'idea' in a sense which is neutral or undecided between the empiricist and the intentional sense, and then ask the question 'What is involved in an idea's being an idea *of* something?' If we then take perception as the paradigm context for ideas, because that is where they originate and is, therefore, in some sense, the source of their credentials, and raise the question of whether perceptual ideas are of something, this can look like a sceptical question. That is, the suggestion that they may not be really ideas *of* something can look to be equivalent to wondering whether there is anything in the world which answers to them; whether, that is, they have the appropriate cause. If, like Locke, one has no patience with scepticism, then one can respond that common sense tells one that perceptual ideas are of something—to think otherwise is mere philosophical scepticism. Having agreed that ideas in their perceptual origin are of something, because they are caused by those things, then whenever they are reactivated in memory and thought, being the same ideas, one will naturally accept that they are of the same things. The conflation of intentionality with epistemological issues is not difficult to understand because both intentionality and knowledge concern ways in which thought makes connection with its referent: but intentionality is logically prior to knowledge of the physical world, because belief enters into the definition of knowledge, and belief is an essentially intentional notion: that is, unless we have some sort of prior conception of what it is for a content to be *putatively* of something, the question of whether it is *in fact* of something cannot be raised. An appropriate kind of causal relation might be what it is for the content to succeed in actually being of something, but that does not explain what it is for a content to *seem to be* of something; that is, it is not an account of intentionality.¹⁶

If this reconstruction is correct, then Locke would have seen no reason to distinguish between ideas in the intentional and empiricist senses. This would facilitate the slide into treating them in the phenomenologically more intelligible empiricist way. Once this conception had gained a grip on the fashionable mind, then philosophers like Berkeley and Hume who are alive to the sceptical challenge are left without a theory of intentionality and are committed to a framework which rules out all but the most reductive attempts at

one. Any theory operating only with the materials that Berkeley and Hume permitted themselves must be reductive; for they have to explain intentionality in terms of some kind of relation between sense-data. It is this predicament that gave rise to associationist accounts of psychology and meaning. The significance of ideas and the flow of thought are no more than the tendency of certain ideas to be followed by others, and ideas cannot be about anything but other elements in the stream of consciousness.

The discussion of Locke's concept of an idea is, however, in a sense, a digression. The adoption of an imagist theory of mental contents in general is a sufficient condition for an empiricist conception of sense-contents, but it is not a necessary condition. Commentators are prone to talk as if holding to an intentional notion of ideas in general saves one from the 'mistake' of treating sensory contents as sense-data.¹⁷ This is not so, for they might have retained a notion of intentionality and yet it may not apply to sensation or perception in the way required; it might only be a property of thought, not of sensation. It is the purpose of the next sections of this chapter to show that there were for progressive seventeenth- and eighteenth-century thinkers strong and intuitively reasonable pressures towards the empiricist conception of sense content that do not rest on a general flight from intentionality, but on reasons that apply specifically to sense-experience.

5 Reasons for the disappearance of the intentional interpretation

The first reason

Almost no one disputes that Berkeley and Hume adopt the full empiricist conception.¹⁸ Nor is it disputed that almost all eighteenth-century commentators and philosophers interpreted Locke and the vocabulary of 'ideas' in that sense. Why did the intentional understanding of ideas die out? More strongly, why did there seem to be such unanimity in reading 'the way of ideas' in the empiricist way, whether or not it was so intended by Locke? There are at least three good reasons.

The simplest reason why Locke and Descartes were read as adopting the empiricist conception is that it seems to follow from their accounts of the physical world. Descartes denies that there is any similarity

between the physical world and ideas, on the grounds that matter possesses only geometrical properties.¹⁹ Locke allows a longer list of primary qualities to the world, but follows Boyle and the atomists in treating secondary qualities as the creations of sense. In itself, the physical world is colourless, odourless, tasteless and silent. This immediately suggests the empiricist conception, for if the secondary qualities are not properties of the physical world, whilst being some of the most obvious objects of our awareness, they must be properties of some non-physical object or content of experience.

It is worth noting in passing that Descartes' conception of matter as possessing only geometrical properties and especially not secondary qualities is inconsistent with that interpretation which attributes to him the scholastic theory that phantasms are physical, for a phantasm that lacks secondary qualities is hardly fitted for the task of endowing perceptual cognition with its distinctive sensory feel—which is the purpose of the phantasm.²⁰ Whether this is an inconsistency in Descartes or a mistake in interpretation I cannot say, but it does mean that the most economic reconstruction of his system is to treat the physical phantasm not as a component of consciousness, for if it were it would have to possess secondary qualities, but rather to treat it as part of the physical structure which, Descartes says, the mind interprets in order to produce conscious states. This is the best reconstruction, that is, if one wishes to maintain his denial of secondary qualities to matter. And this doctrine was so fundamental for the scientific picture of the physical world and, therefore, of the 'modern' philosophy, that it is certainly a feature that a *bien-pensant* late seventeenth- or eighteenth-century reader would wish to retain. Only a thoroughly conservative Aristotelian, who preferred to hold to the qualitative and non-atomist conception of matter and science, would both defy the current atomist wisdom and hold that phantasms are both physical and elements in consciousness and possess a proper range of sensible qualities.

My claim is that it is their exclusion of most sensible qualities from physical reality that makes the empiricist reading of sensory ideas sufficiently natural and reasonable as to explain why it became so quickly the standard interpretation. I shall now consider three objections to this claim; that is, three reasons for not moving from this doctrine of sensible qualities to the empiricist doctrine of ideas. For the present I am only concerned to establish that these three objections do not show the development to be unnatural: to establish

that none of them shows it to be ultimately mistaken will be a task of Chapter III.

1 It might be argued that Locke does not deny secondary qualities to matter: in fact he says that they are a power in matter to produce the corresponding ideas. My interpretation makes them a property of the ideas, and therefore misrepresents Locke. Locke's statement that secondary qualities are powers generally fails to influence his readers because it seems intuitively obvious that phenomenal red ('red-as-it-appears') is not simply in any natural sense a power, and phenomenal red is what we mean by 'red'. It is far more natural to understand Locke's words as meaning that in so far as colour etc. can be said to be in the external physical world at all, it is a power to produce an idea, so the phenomenal red—that which is bright and vivid—is in the idea.

2 Descartes and Locke both argue that sensations and secondary qualities are unclear or indistinct things. Phenomenal red is thus just a confused way of perceiving texture or light on a surface.²¹ This is a doctrine of secondary qualities popular amongst modern materialists and I shall discuss it in more detail later. Suffice it for the moment to say that it is counter-intuitive to treat the phenomenon of colour as nothing more than an uncertain judgement about some more scientifically respectable structure. It appears to be a very plain and obvious phenomenon—much more so than conceptions of microscopic events in terms of which they are to be analysed. It is not surprising that this did not present itself as the natural reading.

3 It will be argued that the fact that certain sensible qualities are sense-dependent does not mean that they are properties of some object of awareness other than the physical world. They are, rather, features of the activity of perceiving. This so-called 'adverbial' approach to the subjective dimension of perception will be considered extensively in Chapter VII. For the moment, all that I require is agreement that secondary qualities—particularly colour—are naturally taken objectively. It seems to me that the experiences of colour which are at least one of the most striking features of space as it strikes us visually, are paradigms of what are, phenomenologically, objects of sensory awareness. If phenomenal colour is best understood, in the end, as non-objectual, that is a fact which needs demonstrating by philosophical analysis. The natural and *prima facie* reasonable understanding of phenomenal colour is that it is an object of awareness. If the world beyond sense does not possess such

properties, it is not surprising or obviously false that readers of Descartes and Locke should almost automatically interpret ideas as the genuine bearers of such properties.

We can conclude that, whatever the force of these three considerations in making it *rationally* unnecessary to move from the empiricist conception of secondary qualities to the traditional conception of sense-data, none of them tends to show that such a move is not overwhelmingly *natural*.

The second reason

The ‘modern philosophy’ arising from Descartes led to a new awareness of the possibility of scepticism—particularly scepticism concerning the external world. A great motor behind such thoughts was Descartes’ postulation of the all-deceiving evil demon.²² The naive interpretation of such thought experiments is that because I cannot prove to myself that my experience is not wholly hallucinatory and subjective then it must be of a sort which is neutral between being physical and being hallucinatory. The only sort of entity which would fit this bill is an idea construed according to the empiricist conception. Such an entity would be an hallucination if it lacked the appropriate cause and a genuine perception if the correct cause were responsible for it.

As it stands, this argument is invalid and not Descartes’. The fact that one does not know whether the contents of experience are real or imaginary does not of itself entail that those contents are something neutral between, or common to, the two. They may be one or the other and one may merely not know which. It is a mistake to think that epistemic possibility concerning the status of an object (that is, cognitive uncertainty about its status) does not allow a real metaphysical possibility that the thing itself might be one thing or another. If, for example, the contents of experience are in fact physical it would be impossible that those very things be nonphysical. Nevertheless, a thought experiment superficially similar to Descartes’ can be conceived which does have the consequence of supporting the empiricist conception of ideas. The assumption is that hallucinations are—or can be—produced by activating the same physiological resources as are involved in perception. These physiological features might be conceived of as phantasms or as animal spirits, or as the nerves and neurons of modern physiology. The point is that the

uniformity of relations between brain and mind guarantee that there is no essential difference between the products in the cases of hallucination and perception.²³ In the case of hallucination the product is some private image and so it must be so in the other case. Like the argument from secondary qualities, this argument will seem obvious to someone accepting the contemporary mechanistic physics, provided they also accept that mind-body interactions are lawgoverned, as Descartes and Locke both did. If the same brain process is always interpreted by the mind in the same way, irrespective of its cause, then how can the idea that the mind forms in this way fail to be neutral between hallucination and perception? How strong this argument really is we shall be considering later, but it is certainly the natural picture to accompany the new science.

The third reason

It would appeal to many modern philosophers to treat the neutral element common to perception and hallucination as somehow intrinsically intentional. David Smith explains the development of the empiricist conception in Berkeley from Descartes' intentionalism in terms of the failure of philosophers to make any clear sense of the notion that an idea might sensorily represent a physical state without itself possessing any salient intrinsic properties.²⁴ The representations that we recognise in other contexts—for example, words and pictures—do not represent intrinsically but by convention or by resemblance. There is a great mystery about how an idea could represent in any other way. Perhaps pure thought might be characterised as an 'act of pure intentionality', but sensory representations, with their manifest sensuous content are not pure in this way: whence comes the phenomenal content we find in hallucinations and 'illusions'. Yolton cites the fact that Locke and Berkeley insistently call ideas of all sorts 'perceptions' as evidence that they subscribe to the *intentional act* rather than the *object* interpretation of them, for a perception is an *act*, not the object of an act. But, even putting aside Locke's description of them as *objects* of perceptions, the assimilation to perception weakens rather than strengthens the intentional case. For, as Locke's own language shows, a perception carries with it its object in a much more literal sense than any other mental act. My pure thought of a lion has an object, but that object does not present itself as literally before me: but in a perceptual state—real or hallucinatory—the object of the act

occupies an area in my sense field or fields. A proper discussion of intentionality must await Chapter VII, but the natural drift of emphasis on perception in our interpretation of ideas is the move towards a ‘picturing’ conception of representation. Even for Descartes, the scholastics and Aristotle, the specifically sensory component in perceptual cognition comes from a non-intentional feature, namely the phantasm. Such a feature can be seen as or projected into the physical world, but in itself operates by resemblance. The increasing and natural emphasis on the sensory paradigm in the understanding of ideas moves one nearer to an empiricist conception of all features of mental life, complete with the imagist theory of thought, associationism and the like. Once the general concept of idea is undermined in this way it is not surprising that, in the case of perception itself, where intentionality is less easy to grasp than picturing or resemblance, the empiricist conception should become the accepted interpretation of mental content as a whole.

6 Intentionality from Reid to Husserl

Reid

I have presented reasons for thinking that the ‘ideas’ of the seventeenth-century philosophers would naturally tend to be understood according to the classical conception; and I have given some reasons for thinking that, though there is a certain opacity in Locke’s mode of expression, he does subscribe to a sense-datum conception of ideas. It is not normal to doubt that by Berkeley—with his anti-abstractionism and imagist theory of thought—the classical conception was firmly established, and intentionality had disappeared as an intrinsic property, not only of perceptual states, but of all mental contents. This has recently been challenged by one of the most distinguished historians of modern philosophy.

According to John Yolton, the philosopher who finally foisted the traditional conception on Locke, Berkeley and Hume was Thomas Reid. This is ironic because Reid is often represented by modern opponents of the empiricists as the outstanding protagonist of direct or naive realism and common sense in the eighteenth century. So, according to Yolton, Reid is able to make himself stand out as a defender of common sense and the intentionality of ideas by misrepresenting the role of intentionality in the writings of his predecessors. I think I have said

enough to show that this is wrong about Berkeley and Hume and, almost certainly, Locke. But Reid's own account of intentionality needs examination just because he is so often invoked by modern opponents of empiricism.

Reid does, indeed, espouse a form of intentionality, but not—as is coming more to be recognised—in a way that has much philosophical depth.²⁵ He differs from the classical conception by denying that ideas or appearances (he accepts both terms) actually possess sensible qualities, and he imports a form of intentionality by claiming that an idea is 'a sign of something external'.²⁶ His account, however, fails to illuminate in a variety of ways. Whilst strenuously denying that colour is a property of ideas, he claims that it is the *unknown* cause of the ideas which we know.²⁷ So the immediately recognisable phenomenon is a property of an idea and the quality itself is knowable only indirectly via its effect—i.e. the appearance.

The name of *colour* belongs indeed to the cause only and not to the effect. But as the cause is unknown, we can form no distinct idea of it, but by its relation to the known effect.²⁸

Locke's mistake is, according to Reid, a purely verbal one, for, once having distinguished the appearance to the eye and the 'modifications of the coloured body' which caused the appearance, the question was 'Whether to give the name *colour* to the cause, or to the effect?'²⁹ If Locke had said, as some say he does (see above), that it was the unknown cause which was the colour, then no problems would have arisen, according to Reid. I have already considered the implausibility of detaching colour words from the phenomenon, but, irrespective of that, it is difficult to see how such a verbal manoeuvre could assist our knowledge or understanding. He has no useful or substantial account of the intentionality of ideas.

By the constitution of our nature, we are led to conceive this idea as a sign of something external... In particular, that idea which we have called the *appearance of colour*, suggests the conception and belief of some unknown quality in the body, which occasions the idea, and it is to this quality, and not to the idea, that we give the name of *colour*.³⁰

The Classical Empiricist Conception

The first objection to this account of ideas as signs is that it is manifestly false that 'by the constitution of our nature' the bright red of a patch suggests to us that 'red' is really the name of an unknown cause of the vivid phenomenon, or that there is such a cause. Our instinct is to take the vivid phenomenon to be the external reality itself, or to project it onto the world. Reid forces together the naive realist understanding of our discourse about sensible properties, with a Lockean epistemology, and, by invoking the unexplained notion that ideas are signs, passes this combination off as pre-theoretic common sense.

Reid accepts the empiricist dichotomy between appearance and reality, even in the much challenged case of perspective:

But yet it is certain that at the distance of one foot, its visible length and breadth is about ten times as great as at the distance of ten feet, and consequently its surface is about a hundred times as great.³¹

Nothing could tend to suggest sense-data more than talk of the 'surfaces' of 'visible appearances'. Reid's only serious point is that we spontaneously adjust or interpret our experiences so as to *see them as* physical objects: it needs training to attend to the phenomena themselves.³² It is this interpretation which constitutes ideas as signs. Exactly how the realist conceptual framework is related to the contents of experience is an interesting and difficult question. Hume thought that the association of ideas could lead to our construction of the physical world from the phenomenal and, hence, the construction of physical concepts. Kant argued that the basic realist categories of space, time, substance and causation had to be known *a priori*. But neither Hume nor Kant were common-sense realists, and the fact that we operate, from as early as we can remember, from the realist perspective is not a strong argument for the truth of direct realism in perception. At the best, an illuminating account of how appearances are signs is required which rules out the view that we merely interpret or project them realistically. Reid does not merely fail to provide such an account, his confinement of the phenomenal and known element in secondary qualities to the level of ideas seems to rule it out.

Brentano

One might hope that the strategy of those opposed to the empiricist conception might become clearer with the explicit resurrection of the

scholastic concept of intentionality by Brentano in the nineteenth century. Brentano, however, is no direct help to the opponents of sense-data, because his view of secondary qualities is fundamentally Lockean.

Knowledge, joy and desire really exist. Colour, sound and warmth have only a phenomenal and intentional existence.³³

He rejects Berkeleian arguments for the sense-dependence of all physical properties, but thinks that there are scientific reasons for affirming it of secondary qualities.

It is not correct, therefore, to say that the assumption that there exists a physical phenomenon outside the mind which is just as real as those which we find intentionally within us, implies a contradiction. It is only that, when we compare one with the other we discover conflicts which clearly show that no real existence corresponds to the intentional existence in this case.³⁴

It might be argued that this is just an accidental feature of his general theory and one which arises from a mistaken understanding of secondary qualities. This reply misses the point. Brentano both clearly thinks that the intentional objects of perception are such that they could be the sole *loci* of colour, sound, warmth, etc. and that such qualities do not exist in the external world. By contrast, what one might call the 'pure doctrine' of the intentionality of consciousness would hold that the contents of consciousness are in all respects *of* something outside consciousness; it would be, as Ginnane says of thought, 'an act of pure intentionality' with no intrinsic, as opposed to intentional, features.³⁵ The location of colour in the intentional object prevents this picture, for the datum is not *of* colour in the external world, for phenomenal colour does not exist in the external world. (I have already considered in connection with Reid and dismissed the suggestion that phenomenal colour is just an appearance *of* unknown colour.)

Brentano is not, however, a pure Lockean, for, though secondary qualities exist only in the senses, they exist there only intentionally, not actually. He argued against treating sense-contents as objects of awareness, possessing real sensible qualities, as follows:

If we investigate what it means to say the colour is not known as actually existing, but as phenomenally existing, it becomes clear

that in the final analysis I do not know that a colour exists, but that I have a presentation of the colour that I see. When I do this I do not mean that the colour exists, *because otherwise someone might suggest that something to which I have a different mental relation, for example, something impossible which I reject as impossible, exists because it exists within me as something denied*. Hence it follows that we do not really recognise that which is known as the ‘object’, what we recognise is only the mentally active being who has it as his object.³⁶

Brentano’s argument is that if we treat phenomenal colour as an object of awareness—something that really is red—we would equally have to treat the round square as an object that really is round and square, which is obviously absurd. This argument entirely fails to allow for the difference between the purely conceptual, where it is not as if an instance of the content is present, and the perceptual. It is plausible to suggest that the intentional objects of most mental attitudes—for example, belief, fear, hope, etc.—do not differ in themselves, and that the differences between these mental states consist of an attitude which is distinct from the object. Thus, if Jones fears meeting a lion and Smith desires to meet a lion the *content* of the attitudes—that they meet a lion—is exactly the same, but Jones has certain feelings about it and Smith has others: the differences of attitudes does not seem to relate to the manner in which the object is in their thoughts, only in their attitude to that object. Perception—called ‘presentation’ by Brentano—is, however, different. If someone actually experiences meeting a lion then this—in abstraction from any feelings they may have about the encounter—involves a different manner of presence for the object itself from all other mental attitudes. It cannot plausibly be said that perceiving the approaching lion differs from thinking of its approaching only in one’s attitude to that content: it differs in the way the content is involved.³⁷ Indeed, thought is essentially a relation to objects that can obtain in their acknowledged absence, whilst perception involves their apparent presence. As intentionality is the defining feature of thought and this connection with possible absence is, therefore, an essential property of it, it would be strange if it could cope with the apparent presence that characterises experience. The presentational nature of experience is why, even if one is convinced that colour is strictly a property that exists only relative to experience, it still seems natural *to see it as* belonging directly to the lion. We cannot

‘paint onto’ the objects of our hopes and fears the subjective content of our attitudes in the same direct way we can paint the subjective aspect of vision onto its objects. This fact is the very same fact as the fact that the subjective feel of perception seems to be not a form of response to an object but a manner of presence of the object itself. For this reason we should be suspicious of the claim that intentionality can cope with the role of the object in perception as completely as it can for other attitudes.

It might be argued that the fact that perception is presentational is irrelevant to its intentionality; the perceptual can be incoherent and is, hence, as intentional and unreal as the conceptual. An Escher drawing, in which a staircase ascends yet joins up with itself, might be cited as an example of such incoherence. But this example works against itself. There is here a perfectly coherent two-dimensional drawing which represents something incoherent as a three-dimensional whole: conceptual, representational properties of a picture can be incoherent, but its basic ones cannot. Given the way that presentation seems to involve actualised secondary qualities, this seems to be a good model for sensations and one that fits well with the sense-datum theory.

Brentano’s theory does, however, exhibit a feature we shall re-discover below, namely the surprising view that nothing actually is coloured. Secondary qualities belong neither to external objects nor, except intentionally, to sensations. This might have some slight plausibility as an error theory: this would say that we originally take colour to belong to external objects, but are forced by science to recognise that it exists only in our experience, and so does not truly exist at all. But if one regards the sense-dependence of secondary qualities as part of our normal understanding of them, one would be in the bizarre position of taking it to be part of our ordinary conception of such qualities that they are not actually instantiated—that nothing could be, for example, coloured.

The move from Brentano to Husserl

The invocation of intentionality to avoid the problems of empiricism in fact faces two problems. The one we have been discussing is that secondary qualities seem to ‘get behind’ intentional objects and be properties of sense-contents, rather in the way that the basic, physical properties of a picture are more fundamental than its representational properties. Given the central place of secondary qualities in the nature

of experience—they are, in a sense, its stuff and guts—then one seems to have, at best, sense-data with some intentional features. (We shall find this again when discussing the percept theory on p. 29.) But there is a more generic problem for the intentionalist approach that has nothing to do with secondary qualities. This is the problem of explaining the notions of intentionality, of ‘intentional object’ and of the supposed ‘intentional existence’ of those objects. I speak of these notions in the plural, but these are really one problem, namely that of explaining what sort of things intentional contents are. The point of Brentano’s strategy is to explain the content of mental acts by reference to their objects. The object is also the content by being an intentional object and intentionally inexistent. But problems arise in trying to understand this. On the one hand, Brentano does not want the intentional object to be a mental entity, like an idea, for the object is what one thinks *about* and that is the thing itself (e.g. the moon) and not an idea of the thing. So it needs to be something extra-mental. But then, on the other hand, there is a problem about non-existent objects of thought. Any plausibility that there is in saying that the moon constitutes the content of my thought about it, by taking on an intentional mode of being, is missing in the case of Pegasus or the highest prime: unless, that is, one is prepared, like Meinong, to allow that these objects do, in some sense, exist. Brentano does not wish to allow some kind of existence to non-existent objects. Brentano never solved the problem of how to avoid either subjectifying the objective or reifying the non-existent.

Husserl tries to avoid this problem by invoking a third element: as well as the subject’s act and its object there is a special kind of structure to the activity called the *noema*. It is this which contains—or is—the intentional content of the act. It has been compared to Frege’s *sense*.³⁸ Fregean sense is the mode of presentation of a reference; it is, that is, a way of thinking about something, a description under which it is apprehended. Similarly, the noema is the aspect or characterisation under which an object is thought. I want briefly to consider three problems with the concept of noemata, two general and one specifically related to its application to perception.

First, the postulation of noemata has the feeling of theft over honest toil. The intentionalist approach to mentality started from a problem and a datum. The problem was to answer the question ‘how can mental states be *about* things?’, and the datum was that when the mind is directed upon something then that object has certain properties that are

summed up in its intentionality. This datum was supposed to constitute the starting point for solving the problem. One way of trying to use the datum to answer the problem is to say that a thought is about an object when the object takes on a special kind of mental existence. This was Brentano's line. Perhaps this approach makes more sense in an Aristotelian framework, where one can say that it is the *form* of the object that is able to have different kinds of realisation. This seems less opaque because we are used to the idea that a universal, F-ness, has two kinds of realisation, one in a thing, thereby making that thing F, and one in an intellect, so constituting a concept of F-ness. Nevertheless, as we saw earlier, on p.10, when discussing classical theories, this strategy fails as an account of the relation between particular objects of mental acts and the acts themselves.

Husserl's solution seems to be simply to say that a mental act is able to be about something and thereby make an object its intentional object by having the *sui generis* property of being able to intend it. It is a 'dormative power' answer, because it simply says that minds mean objects by having states that are able to mean them, and calling these states 'noemata'. Nevertheless, Husserl could be right, in the sense that it may very well be the case that intentionality is an unanalysable property possessed by mental states, and that, in particular, it cannot be analysed in terms of the nature, status or 'activity' of the object, as Brentano hoped.

The second problem is that postulating a *sui generis* intentional state still leaves major difficulties. Mental acts are about things in a variety of ways. The thought that all men are mortal is, in some sense, about men in general: the thought that the Queen of England is wise is about something in a clearer sense, because it is about a particular existing thing: the thought that the mother of the Prince of Wales is wise is about the same thing, but grasped under a different description: and the thought that the present King of France is bald is about a particular, but one with the embarrassing property of non-existence. Because the first thought is quantificational and not referential, we do not worry about the nature of its object and how it enters into the thought. The idea that propositional thought of this quantified kind—purely descriptive thought—is *sui generis* is not too problematic. The difficult issue is how thoughts can refer to particulars, apprehending them under different aspects, especially when they do not exist. This is the problem dealt with in the analytical theory of reference, the territory of Mill, Frege, Russell, Strawson, Kripke, Evans and many

other philosophers. All the thoughts I cite have a descriptive component and the question is whether they have anything else. Russell believed that you could solve the problem of non-existence only by treating them all as pure definite descriptions. Strawson, Evans, (and some say Frege) argue that these thoughts are referential in a way that exceeds their descriptive exactness, so that if they fail of reference there is no real thought at all. Husserl wishes them to be referential in this irreducible way but believes that reference failure does not damage the thought: reference is, for him, an intentional act in a way it is for none of the major figures in the analytic discussion. But his attempt to explain this seems pathetic. In addition to the descriptive component in the noema there is also what he calls ‘the determinable X’, which stands for the object.³⁹ He says nothing illuminating about this. The use of what looks like a variable here must raise the suspicion that he knew he needed something rather like a name or variable, without being able to provide an account of how such an invocation of an object was supposed to work in a way which would give him what he wanted. As an account of intentional objects, nothing has been achieved.

The third problem concerns perception. This is not an objection to Husserl’s theory so much as an objection to invoking it to overcome empiricist problems. The noema is an abstract object, like a Fregean sense; in itself it has no sensory qualities. The noema is the same whether I hope, fear, believe or *see* that the lion is approaching. The qualities of experience are provided by what Husserl calls the *matter* of the experience, which is, roughly, its sensational properties.⁴⁰ The total package looks very similar to sense-data or sensations infused with a physical object interpretation. And this looks like the very moderate modification to the sense-datum theory I mentioned at the outset. To avoid this conclusion we would have to show at least that the sensational component is so intermixed with conceptualisation as physical, that the two cannot be separated. This is what is attempted by the percept theory.

7 The percept theory as an alternative to the classical conception

There is a conception of sense-contents which is an alternative to the classical account, which has been called the ‘percept theory’. It is, in effect, the intentional theory in the form in which it first entered Anglo-

American philosophy. The percept theory shares the first three conditions in my definition of the sense-datum theory, but rejects the fifth and, perhaps, the fourth. So according to the percept theory there is something private and non-physical of which we are aware in perception, but this is intrinsically intentional, in the sense that it is characterised by physical concepts, and, on some versions of the theory, does not possess any sensible qualities intrinsically. Our sense-contents, that is, though distinct from the external world, are essentially structured in terms of physical concepts: we seem to see tables and trees, which seem to have certain colours, we do not sense patches of colour, which we interpret as being, or as being caused by, tables and trees.

Firth claims that, once empiricists allow the reality of depth perception, they are well on the way to the percept theory.

The important fact is that this shift of opinion [about depth] represents a first step towards the recognition that in perception we are conscious of *many* qualities and relations which do not differ in their phenomenological status from those few which have traditionally been attributed to sense-data. Thus it is but one small step, as the Gestalt psychologists have shown, to the recognition that such qualities as simplicity, regularity, harmoniousness, clumsiness, gracefulness...can also have the same phenomenological status as colour and shape. And it is but one small additional step from this to the recognition that the same holds true of qualities fittingly described by such adjectives as 'reptilian', 'feline', 'ethereal', 'substantial' and perhaps most of the adjectives in the dictionary. And this, of course, finally forces the admission that the qualities belonging to objects of direct awareness cannot be thought of as limited, in the manner traditionally assumed, by the use of one or other particular organ of sense.⁴¹

The final step in the Percept Theory consists in showing that the qualities of which we are conscious in perception are almost always presented to us, in some obvious sense, as the qualities of *physical objects*. We are not conscious of liquidity, coldness, and solidity, but of the liquidity of water, the coldness of ice, and the solidity of rocks.⁴²

The contrasting theory, according to which experience consists of a sensory core plus interpretation, is well expressed by Price.

When I say 'This table appears brown to me' it is quite plain that I am acquainted with an actual instance of brownness... But I am not acquainted with an actual instance of tableness, though of course it may be that there is one. Thus the natural way of restating the original sentence 'This table appears brown to me' is 'I am acquainted with something which *actually is brown* (viz. a sense-datum) and I believe there is a table to which this something is intimately related (viz. belongs to)'.⁴³

How are we to choose between these theories? A full discussion of this question will have to wait for later chapters, but it is relevant to setting up our problem to see the *prima facie* reasons either way.

The percept theory has a serious *prima facie* disadvantage, for whilst it is clear what properties sense-data possess intrinsically, it is obscure what—if any—percepts possess. A sense-datum is supposed to be actually, for example, red and square. Something mental cannot literally be reptilian or feline, for to be these things literally is to be a reptile or a cat. It will not do to take these expressions to mean 'reptilelike' or 'cat-like', for we can then raise the question of *how* they are like the real thing, and the candidate answers—in colour, shape or smell, for example—will push us back towards the sense-datum theory. Saying that perceptual consciousness can be feline is saying that it can be *as of a cat* in a way that is not reducible to there being a cat-shaped patch which is interpreted as or seen as a cat. If this is applied generally then a percept of red will not be red but will just be *of an ostensible red object*. Percepts will be one hundred per cent intentional, and this seems less intelligible than the mixture of actual qualities and interpretation that the sense-datum theory provides.

We have already seen in discussing Brentano and Husserl that there are almost irresistible pressures to confine secondary qualities to sense-contents. And Firth admits that secondary-quality terms might be special in applying to percepts in the same sense as they are naively applied to physical objects. Certainly, if there are good reasons for thinking that physical objects are not literally coloured, and one also refuses to attribute them to sense-contents, then one will have the bizarre theory (which, we shall see, has been recently adopted) that *nothing* is actually coloured and that colour exists only in an

intentionally inexistent mode. The desire to avoid such an implausible theory of secondary qualities might force the percept theorist to locate them as properties of sensations and so produce a slightly weaker version of the percept theory. This raises problems about the relation between these sensational properties and the others—often referred to as the *representational* properties of the percept. A colour must have a shape, and once one has admitted to colour patches in sense-contents, it is natural to treat their other representational properties as the products of taking, interpreting or projecting those patches, as the sense-datum theory prescribes. On the other side, there seems to be quite a deep obscurity about what this taking, projecting or interpreting is supposed exactly to be, given that there is no *conscious* process of interpretation, and the fact that experience seems to be of objects and not just patches seems to be fairly basic to the phenomenology.

At first sight, therefore, the sense-datum approach seems preferable, because seemingly less obscure, but the relation between the basic data and the way we construct it requires further investigation. We might discover that there are reasons for attributing intentional properties to sense contents, as well as intrinsic sensational ones, and not putting the former down to activities of the subject.

8 Conclusion

In this chapter I have tried to sketch historically the tension between the emphasis on the pictorial nature of sense-contents, which leads naturally to the sense-datum theory, and the emphasis on their representational role, which tends towards the intentional theory. The almost universal tendency towards treating secondary qualities as subjective puts heavy pressure on the intentional theory and favours the sense-datum interpretation. In the next two chapters I shall consider the classic arguments for adopting the sense-datum theory.

CHAPTER II

The Traditional Arguments for the Empiricist Conception of Sense-contents: the Argument from Illusion

1 The basic argument and the Phenomenal Principle

Probably everyone agrees that objects sometimes look (or sound or feel or taste or smell) different from the way they actually are. A wide variety of different types of case are cited as instances of this obvious fact. Some examples are: mountains look purple when they are not and the sky looks blue when there is nothing actually blue there; clearly defined objects look fuzzy to the short-sighted; different lights make objects look different colours, though the objects don't actually change; distance makes objects look the wrong shape; a hot hand and a cold hand feel the same water as being of different temperatures; Muller-Lyre lines look different lengths when they are not; objects appear to be different shapes from different angles—e.g. a round penny from the side looks elliptical; science shows that physical objects are almost entirely different from how they appear.

Common sense tends to accept these facts (or some or most of them) and to treat them as inconsequential, but many philosophers, including most of the empiricist tradition, have thought that they have an important philosophical consequence, namely that we are not directly or ostensibly aware of mind-independent physical objects. This conclusion is reached because, with some or all of these forms of 'illusion' in mind, an argument of the following kind, informally stated, emerges:

In some/many/most/all cases of perception, we are aware of something that possesses different sensible properties from those possessed by the physical object we take ourself to be perceiving. That of which we are aware is, therefore, something other than the object purportedly perceived.

The argument can easily be made formal, for the premise and the conclusion are linked by an application of Leibnitz's Law: if things possess different properties then they cannot be the same thing.

Two problems with this argument will have to be considered. One concerns its scope. If things do not look other than they are in all cases, then should not the conclusion be similarly restricted, saying that *in those cases only* we are aware of something other than the object perceived? Some philosophers—J.L. Austin, for example—have disputed whether what were supposed to be the major cases of 'illusion' really were instances of objects looking other than they are:

[I]t is simply not true that seeing a bright green after-image against a white wall is exactly like seeing a bright green patch actually on the wall; or that seeing a white wall through blue spectacles is exactly like seeing a blue wall; or that seeing pink rats in D.T.s is exactly like seeing pink rats; or...that seeing a stick refracted in water is like seeing a bent stick.¹

At the other extreme, those who think that science proves that objects are radically different from the way they appear think that every case of perception is a case of things looking other than they are. It is necessary, therefore, to sort out two things:

- 1 Which, if any, of the purported cases of 'illusion' really are cases of things appearing other than they are.
- 2 Assuming that such mis-appearances do not cover all perception, whether there is any way of generalising the conclusion.

The second problem is more important, because less easily solved. It concerns the relation of intentionality and perception—a theme on which we have already touched and to which we shall repeatedly return. In a sense, the major issue of this book is whether intentionality can provide the basis for an alternative to the classic conception of sense-contents. The premise of the argument from illusion rests on what I shall call the *Phenomenal Principle* (=P):

(P) If there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality.

Much of the discussion in this chapter from p. 37 will concern this principle, but a final decision on it must await Chapter VII.

The argument and early modern philosophy

Cases of misperception such as those reported above have often been cited through the history of philosophy, and they figure in the writings of Descartes, Locke, Berkeley and Hume. It is natural, therefore, to read the passages in which these phenomena are cited as presenting a version of the argument from illusion. This would, however, be too swift, and comes from conflating issues in the philosophy of perception with sceptical problems in the theory of knowledge.

As a result of the philosophy of the early modern period, the philosophy of perception and the problem of scepticism became very closely intertwined, but they had not originally been so closely connected. Ancient scepticism operated by trying to prove that whenever we think that we have reason to believe that *p*, we also have an equally strong reason to believe that *not-p*, and so there ought to be a suspension of belief. The fallibility of perception, as exhibited in the sorts of experience cited in the ‘argument from illusion’—that is, experiences in which things look other than they really are—is invoked as one side in one kind of clash of evidence.² But there never is an argument of the form ‘what we are aware of in perception is some kind of mind-dependent representation, not an external object, so we have no good reason for thinking (or, at least, have a problem about showing) that there is anything existing outside our minds’. It is in this form—through the reification of appearances as ‘ideas’, ‘representations’, ‘sense-data’ or ‘*qualia*’—that perception and scepticism have become interconnected in the modern era; that is, the connection is via the acknowledgement of a ‘veil of perception’ hiding the world from us. It was, for example, fear of scepticism based upon representative realism that motivated Berkeley’s idealism. This change of emphasis is remarked on and approved of by Hume in the *Enquiries*:

I need not insist upon the more trite topics, employed by the sceptics in all ages, against the evidence of *sense*; such as those that are derived from the imperfection and fallaciousness of our organs, on numberless occasions; the crooked appearance of an oar in water; the various aspects of objects, according to their different distances; the double images which arise from the pressing one eye; with many other appearances of a like nature. These sceptic topics are only sufficient to prove, that the senses alone are not implicitly to be depended on; but that we must

correct their evidence by reason, and by considerations, derived from the nature of the medium, the distance of the object, and the disposition of the organ, in order to render them, within their sphere, the proper *criteria* of truth and falsehood. There are many more profound arguments against the senses which do not admit of so easy a solution.³

The real sceptical threat comes from the fact that we are instinctively committed to a naive realism that tells us that the very contents of our perception are mind-independent, whereas

this universal and primary opinion of all men is soon destroyed by the slightest philosophy, which teaches us that nothing can ever be present to the mind but an image or perception, and that the senses are only the inlets, through which these images are conveyed, without being able to convey any immediate intercourse between the mind and the object.⁴

At this point Hume does employ the argument from illusion to show how easily the 'universal and primary opinion of all men' is destroyed by reflection. His argument is perfunctory, but only because he thinks it as painfully obvious as it is crucial.

The table which we see, seems to diminish, as we remove farther from it: but the real table, which exists independent of us, suffers no alteration: it was, therefore, nothing but its image, which was present to the mind. These are the obvious dictates of reason; and no man, who reflects, ever doubted, that the existences, which we consider, when we say, *this house* or *that tree*, are nothing but perceptions in the mind, and fleeting copies or representations of other existences, which remain uniform and independent.⁵

He then raises the sceptical doubt:

By what argument can it be proved, that the perceptions of the mind must be caused by external objects, entirely different from them, though resembling them (if that be possible) and could not arise either from the energy of the mind itself, or from the suggestion of some invisible spirit, or from some other cause still more unknown to us?⁶

When Descartes, at the start of the *First Meditation*, cites facts such as that distance distorts the shape of objects, he is following the ancient

sceptical strategy by arguing that perceptions conflict with each other and perception is, therefore, unreliable. Like Hume, he thinks this a weak argument, because we can appeal to reason to sort out the reliable from the unreliable cases. He, therefore, moves on to more radical suggestions about dreaming, and, finally, the evil demon. His canvassing of cases of misperception is, therefore, much more similar to the ancient one of which Hume thought little than it is like an argument about representative realism. He does not, however, put much weight on them.

I doubt, too, whether Locke used the argument from illusion. When he cites the relativity of perception—as in the case of water feeling warm to one hand and cold to the other—this is meant to show the sense-dependence of *secondary qualities* only.

Whereas it is impossible, that the same Water, if those *Ideas* were really in it, should at the same time be both Hot and Cold. For if we imagine *Warmth*, as it is *in our hands*, to be *nothing but a certain sort and degree of Motion in the minute Particles of our Nerves, or animal Spirits*, we may understand, how it is possible, that the same Water may at the same time produce the Sensation of Heat in one Hand, and Cold in the other; which yet Figure never does, that never producing the *Idea* of a square by one Hand, which has produced the *Idea* of a Globe by another.⁷

This is a terrible howler. Perceptual ‘illusions’ show—if anything—that sense contents are mind-dependent, but Locke is using this instance to show that the secondary qualities that figure in the illusion exist only in the mind. He claims—by choosing an easy case—that this kind of perceptual relativity does not apply to primary qualities. Berkeley has no difficulty showing that they do so apply, and that, therefore, primary qualities are equally mind-dependent and there are no clothes left for a *mind-independent* world to dress up in.

It might seem surprising that there is no straightforward use of the argument from illusion in the main texts of Descartes, Locke or Berkeley, for we associate this and the other arguments for the classical conception of content with the Cartesian and empiricist traditions of mind and perception. We have to wait for Hume before the argument occurs, but he makes matters more mysterious by referring to it as if it were a commonplace amongst modern philosophers, which would lead us at least to expect it in Locke.

The absence of the argument would be grist to the mill of those who think that Descartes, Locke and even Berkeley lacked the classic conception of sense-contents, for the argument rests on that conception. I explained above that the argument required the Phenomenal Principle, which is the principle that the contents of experience are not intentional, but do instantiate sensible qualities. Those who think that the *ideas* of early modern philosophy are intentional would, therefore, not expect to find the argument from illusion in those sources.

For reasons already explained, I do not think that ideas can be intentional, though there may be some uncertainty about how self-conscious those philosophers were about the issue. However inexplicit they might have been about it, I think there is one very powerful reason for believing that any philosophers in a broadly Cartesian tradition—as are all those we are considering—must have accepted the Phenomenal Principle. This is that the principle is very closely associated with a belief in the self-intimating nature of consciousness. All these philosophers believed that there could be no certainty greater than that given by something's being clear and distinct to consciousness; plain phenomenology is the foundation of everything. And nothing could be more clear and distinct than that if one steadily seems to see something red then one is indeed aware of red. (That one is aware of *red* but not *something red*, is a distinction which, with their nominalist ontology of qualities, the empiricists, at least, could not have made.) Once one takes the Phenomenal Principle for granted, the argument from illusion can seem to have become a truism that is hardly worth stating.

The argument that if the intentionalist interpretation of *ideas* were correct, one would not expect to find the argument from illusion, can be matched by the other side: if the imagist interpretation of ideas were correct, the argument from illusion would not be needed. The argument from illusion is an argument against naive realism. If the traditional interpretation of 'the way of ideas' were correct, and ideas tended to be assimilated to images, then the foundations of the system would presuppose the falsehood of naive realism. If everything the mind concerns itself with is something like a mental image, more or less by definition, then the question of whether it can be directly acquainted with mind-independent objects will not arise, and an argument against it will not be required.

Another possible reason for the argument's absence is the subordination of the philosophy to science. There is no great attempt made

to distinguish and label the various arguments used about the nature of perception and its objects; and this is because they are not thought of as separate because they all follow from and are inspired by the new science: they invoke, that is, supposed facts about atomic causal processes, and how experience depends on such processes, and about the dispensability of secondary qualities from the atomist conception of the physical world. The argument from illusion, on the other hand, is a piece of pure phenomenology, and owes nothing to the new science. It is, therefore, not relevant to the question of how that science should be philosophically received.

The argument and early twentieth-century philosophy

The argument from illusion and its cornerstone, the Phenomenal Principle, come into their own in the early twentieth century, when the phenomena of ‘illusion’ tended to replace scientific facts as the basis for arguments in the philosophy of perception. That this represented a real development was noted by the astutest sense-datum theorist, C.D. Broad:

This type of theory, though it has been much mixed up with irrelevant matter, *and has never been clearly stated and worked out till our own day*, is of respectable antiquity. The doctrine of ‘representative ideas’ is the traditional and highly muddled form of it.⁸

Broad, too, has an explicit statement of a developed version of the Phenomenal Principle:

Whenever I truly judge that x appears to me to have the sensible quality q , what happens is that I am directly aware of a certain object y , which (a) really does have the quality q , and (b) stands in some peculiarly intimate relation, yet to be determined, to x .⁹

Broad calls the object of direct awareness a ‘sensus’ and G.E. Moore calls it a ‘sense-datum’. The question of naive realism then becomes the question of whether a sensum or a sense-datum could be identical to the perceived part or feature (usually thought of as the surface) of a physical object. As the quality q is often different from any quality possessed by the relevant physical object or surface, the answer is

plainly that they are not identical, and that, therefore, what we are directly aware of is something other than the physical object we take ourselves to be perceiving. G.E.Moore makes a classic statement of the argument:

But now let us consider the case in which we are not prepared to assert that the surface in question [i.e. the physical surface we think we are perceiving] has changed perceptibly. The strange fact from which the argument I mean is drawn, is that, in a very large number of such cases, it seems as if it were unmistakably true that the presented object, about which we are making our judgement when we talk of 'This surface' at the later time *is* perceptibly different, from that about which we are making it when we talk of the surface I saw just now. If, at the later time, I am at a sufficiently greater distance from the surface, the presented object which corresponds to it at the time seems to be perceptibly smaller, than the one which corresponded to it before. If I am looking at it from a sufficiently oblique angle, the later presented object often seems to be perceptibly different in shape—a perceptibly flatter ellipse, for instance. If I am looking at it, with blue spectacles on, when formerly I had none, the later presented object seems to be perceptibly different in colour from the earlier one....All this seems to be as plain as can be, and yet it makes absolutely no difference to the fact that of the surface in question we are *not* prepared to judge that it is perceptibly different from what it was....It seems, therefore, to be absolutely impossible that the surface seen at the later time should be identical with the object [i.e. the object of immediate awareness, or sense-datum] presented then, and the surface seen at the earlier identical with the object presented then, for the simple reason that, whereas with regard to the later seen surface I am not prepared to judge that it is in any way perceptibly different from that seen earlier, it seems that with regard to the later sense-datum I cannot fail to judge that it *is* perceptibly different from the earlier one: the fact that they are perceptibly different simply stares me in the face....This is the argument, as well as I can put it, for saying that this object is *not* identical with this part of the surface of this inkstand.¹⁰

Almost immediately, however, he expresses a reservation:

But nevertheless it does not seem to me to be quite conclusive, because it rests on an assumption, which, though it seems to me to have great force, does not seem to me quite certain. The assumption I mean is the assumption that, in such cases as those I have spoken of, the later presented object really is perceptually different from the earlier. This assumption has, if I am not mistaken, seemed to many philosophers to be quite unquestionable; they have never even thought of questioning it; and I own that it used to be so with me. And I am still not sure that I may not be talking sheer nonsense in suggesting that it can be questioned. But, if I am, I'm no longer able to see that I am. What now seems to me to be possible is that the sense-datum which corresponds to a tree, which I am seeing, when I am a mile off, may not really be perceived to *be* smaller than the one, which corresponds to the same tree, when I see it from a distance of only a hundred yards, but that it is only perceived to *seem* smaller; that the sense-datum which corresponds to a penny, which I am seeing obliquely, is not really perceived to *be* different in shape from that which corresponds to the penny, when I am straight in front of it, but is only perceived to *seem* different—that all that is perceived is that the one *seems* elliptical and the other circular; that the sense-datum presented to me when I have the blue spectacles on is not perceived to *be* different in colour from the one presented to me when I have not, but only to *seem* so... If such a view is to be possible, we shall have, of course, to maintain that the kind of experience which I have expressed by saying one *seems* different from the other—'*seems* circular', '*seems* blue', '*seems* coloured', and so on—involves an ultimate, not further analysable, kind of psychological relation, not to be identified either with that involved in being 'perceived' to be so and so, or with that involved in being 'judged' to be so and so; since a presented object might, in this sense, *seem* to be elliptical, *seem* to be blue, etc., when it is neither perceived to be so, nor judged to be so. But there seems to me to be no reason why there should not be such an ultimate relation.¹¹

Moore cannot bring himself to believe the 'seems' analysis, for it seems to be a travesty of the phenomena, but he cannot see how to prove this.

The great objection to such a view seems to me to be the difficulty of believing that I don't actually perceive this sense-datum to *be* red, for instance, and that other to *be* elliptical; that I only perceive in many cases that it *seems* so. I cannot, however, now persuade myself that it is quite clear that I do perceive it to *be* so.¹²

I have quoted Moore's dialogue with himself at length because it expresses the essence of the controversy about the argument from illusion. On the one hand, there is the seemingly obvious fact that when one seems to see an object with a certain sensible quality, such as redness, nothing could be plainer than that one is aware of red. On the other hand, language is well equipped with idioms involving 'seems', 'looks', 'appears', which can be used to describe these situations; and it is a feature of these idioms that things can seem, look or appear so and so without their actually being so and so. Indeed, the essence of these expressions is that they exist by the contrast between *really being* so and so and only being so *apparently*.

Moore himself came later to think that the 'seeming' theory was, indeed, nonsense, as he had suspected, because such idioms could only be understood in terms of sense-datum ones:

And it also seems to me plain that, to say that, e.g. if I am wearing blue spectacles, a wall which is white but *not* bluishwhite 'looks' bluish-white to me, is merely another way of saying that I am directly seeing an expanse which really is of a bluish-white colour, and which at the same time has to the surface which is not bluish-white a specific relation....

If I am *not* directly seeing a bluish-white expanse which *has* some such relation to a wall which is *not* bluish-white, how can I possibly know that that wall *is* looking bluish-white to me? It seems to me quite plain that I cannot 'see' in the common sense any physical object whatever without its 'looking' *somehow* to me, and, therefore, without my directly seeing some entity which has [the appropriate relation] to the object I am said to see.¹³

H.H.Price, in a famous passage, also affirms the obviousness of the Phenomenal Principle:

When I see a tomato there is much that I can doubt. I can doubt whether there is a tomato that I am seeing, and not a cleverly

painted piece of wax. I can doubt whether there is any material thing there at all. Perhaps what I took for a tomato was really a reflection; perhaps I am even the victim of some hallucination. One thing however I cannot doubt: that there exists a red patch of a round and somewhat bulgy shape, standing out from a background of other colour-patches, and having a certain visual depth, and that this whole field of colour is directly present to my consciousness. What the red patch is, whether it is physical or psychical or neither, are questions that we may doubt about. But that something is red and round then and there I cannot doubt.¹⁴

Once one has reached this point, the direct realist is sunk, for, in many cases at least, the qualities one is directly aware of are not qualities possessed by the physical object perceived.

Unease is, however, a natural reaction to this argument. There is something close to a paradox in the sense-datum theorists' conception of naive realism. They say that naive, pre-philosophical common sense holds that we are directly aware of external objects, and that, therefore, the naive realist is committed to the view that objects never appear to be other than they are. This latter seems to be a strange idea to attribute to common sense, however naive. Naive realists might be thought to be making either or both of two mistakes. They might be thought to be failing to notice that objects do often look other than they are: or they might be accused of not realising that perceiving objects directly should entail seeing them as they are. The main point is, I think, the latter.

There is a certain historical irony in this critique of naive realism which is, I think, worth remarking. The sense-datum theorists were major figures in the revolt against Hegelian idealism: the argument of Moore's 'Refutation of idealism', for example, turned on the nature of the contents of perceptual consciousness, and Chapter I of Price's *Perception* is largely directed against the holism of the idealists. But there is a striking similarity in structure between the attack on naive realism and certain classic arguments of the British idealists. The similarity consists in attempting to show that a fundamental commonsensical notion contains a contradiction. McTaggart argued that our ordinary notion of time contained two irreconcilable elements: Bradley thought he could prove that there was a contradiction in the idea of distinct but spatially related objects.¹⁵ Moore, Broad and Price

similarly believe that what they take to be our pre-philosophical idea of what it is to perceive the world contains two incompatible elements. One is that we are in direct perceptual contact with the external world, which is taken to imply that one is connected to the world by a 'transparent' conscious apprehending which simply puts one in touch with the object as it is. The other is that objects vary in appearance although they are not changing intrinsically, so that we are not in touch with the object as it is intrinsically. Pre-philosophical common sense about perception is, therefore, self-destructive.

The way many philosophers have reacted to this refutation of naive realism is similar to the way in which the early analytical philosophers reacted to the Hegelian proofs of the unreality of time or space, namely with the strong intuitive feeling that a conception which is so familiar and fundamental as our perceptual realism cannot really be flawed; the problem must stem from philosophical over-simplification or legerdemain. One way of understanding discussion about the 'language of appearing' is as an attempt to get clear about whether the only intelligible conception of consciousness is of it as transparent—that is, whether it is simply a grasping of its object, such that any *apparent* difference in the object has to be attributed to a *real* difference of object, and not to a different way of being consciously related to the same object.

Winston Barnes, in one of the most influential attacks on the sense-datum theory, takes this to be the argument from illusion's principal weakness.¹⁶ He represents it as a contradiction in the argument. He expresses the argument in the following example. The premises are

- 1 I see the rose.
- 2 The rose appears pink to me.
- 3 The rose is red.

The conclusion is

- 4 I see a pink sensum.

He points out that the conclusion contradicts one of the premises, namely (1). He correctly remarks that the argument only works if one assumes that *really* seeing the rose ought to involve seeing it as it is, and this assumption, he argues, contradicts (2) and (3). He seems to think that these contradictions show the falsehood of the argument, but this is question-begging; they could equally be taken as showing the contradiction in our normal beliefs, which the argument is exposing

and exploiting. The core of Barnes' argument is that (4) simply does not follow from the premises.

I cannot infer, as is proposed, *merely from the three facts that I am seeing something, that it looks pink and that it is red, that there is a pink something where the thing appears pink to me.*¹⁷

Barnes' tone is as if this were a knock-down *logical* point about the validity of an inference; and he is right to imply that there are, indeed, plenty of cases in which something's appearing F when it is not does not imply that something *else* is F instead. If something appears to be a unicorn and isn't, nothing else need be a unicorn either; and if John appears angry, but isn't, I'm not aware of some other angry thing. This is just to say that these idioms take intentional objects, and no one denies that. It is, however, important to distinguish between the logical features of psychological states with intentional objects and their phenomenological ones. When I consciously think something about a unicorn there is no unicorn of which I am thinking; but the conscious episode is not constituted simply by 'absence of unicorn' but by some positive features, including, in this case, saying some words to myself. The issue with illusion does not concern the logic of 'appears' and similar words, but how to account for the phenomenology of appearance. The episode of a red rose's appearing pink to me is no more accounted for by the *absence* of pink than the episode of thinking about a unicorn is accounted for by the absence of unicorn. Furthermore, the psychological reality of appearances cannot be explained by the merely conceptual presence of something, as one might do for pure thought (though even there some sort of vehicle, like language, seems to be involved in conscious thought). The specially sensational or presentational element of perception must be accommodated, which the sense-datum theory does, but merely pointing out the intentionality of 'appears' does not.

The failure to distinguish the phenomenological from the logical which tends to go with emphasising intentionality and assimilating perception to thought leads to equivocation or contradiction in ontology. Barnes says

Modes of appearance are clues to the nature of what exists, not existents. I submit that it is improper to ask whether the pink mode of appearing, which is how the rose appears to me, exists.

Although modes of appearance are not existents, they are the material and the only material on which thinking can operate to discover the nature of existing things.¹⁸

We shall investigate further in Chapter VII attempts to dodge ontological questions about appearances and intentional objects, but, at first sight, it is difficult to see why it should be improper to ask whether something exists and even more difficult to see how something that does not exist could be a ‘clue to the nature of what exists’ or the ‘only material on which thinking can operate to discover the nature of existing things’. These obscurities can be exposed and explained if we see them as resulting from a confusion between the logical properties of certain verbs—which can take non-existents as their objects—and the concrete reality of psychological phenomena, which can never be explained just by reference to that non-existence. This confusion pervades modern discussions of appearance.

2 Alternative accounts of appearance

Multiple location

We can say, with some arbitrariness, that the attempt to understand experience in terms of *appearing* etc. has taken five forms. The first three were discussed by early twentieth-century philosophers, and I shall deal with them here. Numbers four and five will have to await Chapter VII. The five theories of appearing are: (1) the multiple location theory; (2) the relational theory; (3) the judgemental theory; (4) the adverbial theory; (5) the intentional object theory. In some of their versions these theories overlap.

The multiple location theory can be regarded as an attempt to reconcile the sense-datum theory with naive realism, rather than as an alternative to the sense-datum theory. Its purpose, that is, is to explain how the radically different appearances that objects present could, nevertheless, be actual parts or aspects of the object ‘out there’: it is a way of explaining how the various and incompatible appearances—sense-data—are all intrinsic to the object. A modern reader might wonder why a naive realist should have any concern to retain sense-data: he must remember that retaining sense-data, in this context, consists essentially in no more than believing that when one clearly seems to see something, e.g. red, then one really is aware of

(something) red. As our discussion of Reid, Brentano and Husserl above suggested, up to this point no one had yet managed to bring themselves to deny this in more than a verbal sense. We have already seen in the discussion of Barnes how slippery such a denial is. So, although the multiple location theory sounds rather bizarre to us, it deserves serious consideration because of the relatively clear way in which it reconciles phenomenology with direct realism. It is, perhaps, the only version of direct realism which seeks to reconcile the claims of observer relativity with those of objectivity without being shifty about the reality of the phenomena. As such it merits serious consideration. Price states the theory as follows:

The Theory of *Multiple Location* says that we must distinguish between the characteristics which characterize something only *from a place*, and those which characterize it *simpliciter*. From a particular place the penny's top surface really is elliptical and smaller than the top surface of a sixpence: but simply—in itself or from no place—it is circular and twice as large. In this way two sense-data having different shapes and sizes, provided they have them from different places, may both be identical with the same part of a surface of a material object (say with the top of it): the one with the top-from-this place, the other with the top-from-that place.¹⁹

And graphically characterises it:

Thus even the humblest and simplest material object, such as a table or a penny, is really a sort of infinitely various porcupine, which is not merely here in this room (as we commonly take it to be) but sticks out as it were in all sorts of directions and to all sorts of distances, 'from' all of which it has its being and is qualified in various ways, whether present to anyone's senses or not.²⁰

Price, in a discussion which is sometimes opaque, seems to have two connected objections to the theory. The first is that he is 'not at all clear' that 'one and the same entity can *both* be qualified from a place, *and* be qualified simply from no place', or that 'the *same* entity can have at the same time an infinite multitude of qualities from an infinite multitude of places'.²¹ Then he claims that it does not seem to cope with double vision or with hallucinations, for these are not dependent on altering viewpoint on something.²² The first objection seems to be

intuitive and not backed up by argument. The second only shows that this account does not cover all kinds of illusion. This will not be too serious if those kinds it does not touch seem intuitively to be of a different sort from the rest, as is *prima facie* plausible for both double vision and hallucinations, for neither are apparent qualities of some actual object; (doubleness is not, strictly, a quality of a thing and, in hallucination, there is no actual object).

Nevertheless, there are two strong arguments against the multiple location theory. First, the theory does not naturally accommodate variations in appearance that are caused by changes in the perceiver of a much more straightforward kind than pressing an eyeball to see double. If I take off my glasses everything becomes blurred. Are we to assume that there are blurred sensibilia waiting to be perceived by someone not only in my location, but with exactly my degree of short-sightedness? Even if one were to swallow this, there would still be difficulty in seeing how a *blurred* sensibilia could be identical with a *well-defined* surface: the object itself would become blurred. One might hope to have a different account of cases where differences depend on subject and not merely on location. This will not do. It seems reasonable to think that there is probably almost universal variation between people, to some degree, of exactness of vision, which means that the additional theory required to cope with subjective variations would be needed for all perception. But, putting this probability aside, any theory which required a different account of what is going on when I perceive with and without my glasses, or as I get tired, or tired and emotional, and see things less clearly, or in sun-glasses, seeing colours darker, loses all plausibility. It is, I suppose, conceivable that the theory that coped with subjective variation might be combined with a 'multiple location' account of public perspectival variations, but my suspicion is that an account that worked for subjective variations would accommodate perspectives, for the two kinds of case seem continuous.

Second—and this may be Price's intuition—it is difficult to see how an elliptical sense-datum could be identical with a round surface, for anything identical with a round surface would have to be round. It could be identical with an elliptical *part* of the round surface, as if an ellipse had been drawn on the penny, but this would have the severe disadvantage that the outline of the ellipse would not, for the most part, coincide with the edge of the penny, which it should, for the outline of the ellipse is meant to *be* the outline of the penny. Similar problems

will occur for the smaller-than-life sense-datum of the penny seen from a distance. These problems could only be solved whilst continuing to hold that the data are identical with the relevant part of the physical object by saying that the datum is round, but looks elliptical, or full size and appears small. But this is to reinstate the problem that sense-data are meant to solve.

It might be replied that this objection misses the point. I am talking as if being *elliptical from p* involved being *elliptical in itself*, which is explicitly not the case: the same surface really is multiform, but only *from* different viewpoints, and this does not involve its being multiform *at* its own location. This reply will not do. The theory is combining naive realism and sense-data. Because of the latter there must be something which is elliptical and because of the former it must be in physical space—indeed, we have been told that it is identical with the surface of the penny, so that must be its location.

To save itself, the multiple location theory must develop in either of two directions. It could avoid the embarrassment of having to identify something elliptical with something round by denying that anything is elliptical when an object is *elliptical from p*: that is, it could treat this as an unanalysable property. This is essentially a move to the *relational theory*, the core of which is the claim that the property of *appearing F to x* is basic and not to be analysed in terms of anything being F. We shall discuss this next.

The second approach is discussed by Price and, in a sense, defended by Ayer in *Foundations of Empirical Knowledge*.²³ The problem of identifying something elliptical with something round could be avoided by refusing to require that the public elliptical datum seen from a certain angle is identical with the round surface. A physical object will, on this account, possess a cluster of sensible properties, each sensed from a different perspective and by a different type of perceiver, but which do not constitute a geometrically solid object. The object will be more like a bundle of intersecting sensible planes, each accessible only from the right perspective and by the right kind of observer. This saves naive realism about perception by entirely abandoning our ordinary concept of a physical object, which requires it to be, in Price's phrase, *spatially unitary*.²⁴ Not merely does this subvert a part of common sense which is much more fundamental to our conception of the world than is the naive realist approach to perception, it is difficult to see the connection between this concept of body and the scientific one: the common-sense

spatially unified solid seems to be continuous with the scientific conception of matter in a way that a mere concatenation of sensible surfaces is not. It is also probably true to say that the standard conception of physical objects is built into our naive realism, for it affirms our direct perception *of bodies as we normally conceive of them*, not of some fantastic construct.

The relational theory

Price calls this 'the theory of appearing' and this is its more usual name, but I have avoided it here because all five theories are, in a sense, theories of appearing. It takes uncontroversial descriptions of the form 'x appears F to S' and construes them as expressing an unanalysable relation between an object, or surface of an object, and a subject. Price states it clearly:

Now the same top surface of a certain penny stamp may appear to me pink and to a colour-blind man grey, to me lozengeshaped and to him trapeziform, while in itself it is square and (perhaps) colourless. Of course the same entity cannot *be* at once red and grey and colourless, trapeziform, lozenge-shaped and square. But then it does not have to be. For though *being* trapeziform is incompatible with being lozenge-shaped, yet appearing trapeziform to A is perfectly compatible with appearing lozenge-shaped to B—and with being intrinsically square. So, appearing grey to B is compatible both with appearing pink to A and with being intrinsically colourless.²⁵

Price points out that this theory faces the same problems as 'multiple location' when it comes to seeing double and hallucination. But his fundamental objection is that 'its own foundations are incoherent':²⁶

Let us now turn back to its fundamental expression 'A appears *b* to So-and-So'...and ask what meaning we can give to it. Clearly if A is the name of a material object such as a table, it does mean something. But equally clearly, that meaning is further analysable. When I say 'This table appears brown to me' it is quite plain that I am acquainted with an actual instance of brownness (or equally plainly with a pair of instances when I see double). *This cannot indeed be proved, but it is absolutely evident and indubitable.*²⁷

He repeats this conviction later:

We cannot on reflection doubt that there exists an instance of brownness.²⁸

Price is refuting the theory of appearing by taking the phenomenal principle as beyond reasonable doubt. As the phenomenal principle is exactly what is at stake nowadays, Price seems to be doing no better than begging the question. There is food for thought here, however. First, this is another illustration of the fact already noted that, at least with regard to secondary qualities, it is difficult to find philosophers of any persuasion in the past who did not take the phenomenal principle to be an obvious datum. Second, the fact that nowadays its denial is almost a commonplace does not show that Price and these earlier philosophers were wrong: perhaps more recent philosophers were just taken in by bad reasons for denying the obvious, whilst their judgement was clouded by an urgent desire *not* to believe in sense-data. One of the purposes of this book is to show that the rejection of intuition regarding the Phenomenal Principle rests on a tissue of confusions, but no coherent case.

There is a powerful argument against the relational theory as I have described it. It is a semantic one. 'X appears F to s' is presented as being semantically simple. I said earlier that it is presented as a two-term relation between object and perceiver; so 'appears F to' is an unanalysable relation. (This means, of course, that x and s are variables, but F is a name. This might be clearer if 'x appears square to s' was given as an example of the sort of relation concerned.) It would seem that this must be wrong, because the relation contains a predicate which can occur in other contexts as a part: 'F'—or 'square'—is not equivocal between 'is F' (or 'is square') and 'appears F' ('appears square').²⁹

Now, on Price's formulation of the theory, 'F' is a variable and the relation is a three-place one between an object (more exactly, part of the surface of an object), a property and a subject. This avoids the problem that flows from trying to take 'appears square' to be unanalysable, but it faces others.

The problem with this version is that it fails to explain the nature of appearing: it is too generic and, as far as it goes, no one need disagree with it. That is, no one need deny that, when something appears square to someone then there is a relation of some sort involving a thing, a person and the property *square*. The problem is to say something

illuminating about the nature of the relation in question, particularly about how *square* comes into it. The fact that the relation is unanalysable does not free us from the obligation to give some guidance on its nature, for there are other relations that involve object, subject and property (e.g. 's thinks x is F') and there must be something that distinguishes appearing from the others. Unanalysability might mean we cannot *say* what this is, but then the differences will have to be ostensively recognisable. This will lead the theory back into trouble.

We can recognise the difference between thinking a thing to be F and its appearing F; that is, we can recognise the difference between apprehending a property intellectually and apprehending it sensibly. But if the 'appearing' relation is the unanalysable sensible apprehension of a property with respect to an object, where that property need not actually be instantiated in that object, we seem to have something very close to the sense-datum theory, for we are involving the property *in a sensible form*, yet not in the physical object. The only difference from the sense-datum theory is that what we are aware of is a species of universal, not a particular—it is, that is, a *quale* rather than a sense-datum. But *qualia* are close cousins of sense-data, and the question of which of these two ways better captures the nature of sense-contents is very much a domestic dispute within the empiricist family.

So relational accounts which take the relation to be two-term, thus preserving the realist intuition that the observer and the physical situation are the only things involved in perception, make the mistake of suppressing semantic complexity: and those which bring in the property as a third relatum, *and* acknowledge the difference between the intellectual and the sensible presence of a property in a mental state, fall back into the sense-datum theory.

The judgemental theory

According to this theory, when *a* appears F to *s* and *a* is not really F, what is occurring is that *s* is judging, or is inclined to judge, or something similar, that *a* is F. In this way the phenomena that fall under 'illusion' are represented as cases of mistaken belief, or, at least, of entertaining mistaken beliefs, or feeling some sort of inclination to mistaken belief. The point of this theory is that no one nowadays thinks that belief or judgement—mistaken or otherwise—

essentially involves sense-data. Therefore, if 'illusions' are merely belief-phenomena of some sort they cannot require the postulation of sense-data.

The judgemental theory of *illusion* must be distinguished from the judgemental or belief theory of *perception as a whole*.³⁰ According to the latter, veridical perception consists only in the acquisition of belief through the senses. Its motivation is similar to the more restricted judgemental theory, namely to eliminate a troublesome, irreducible notion of experience, but it is generally adopted by physicalists who find *experience as a whole* an embarrassment, rather than by those who need to explain 'illusory' phenomena, and think they can take veridical perception as unproblematic. We shall be discussing this in Chapter V on physicalist theories of perception.

A little reflection shows the restricted judgemental theory to be inconsistent. If the experience of seeming to see something red when nothing relevant is red is adequately analysed in terms of acquiring some kind of belief state, then a similar analysis must also be adequate for the experience of seeming to see something red when that experience is veridical; for, *qua* experiences, the two are indistinguishable.

C.D.Broad pointed out many years ago that the sense of 'seems', 'looks' or 'appears' relevant to 'illusions' is not the judgemental one. He recognised that this point was 'certainly of the utmost importance if it be true.'³¹

Appearance is *not* merely mistaken *judgment* about physical objects. When I judge that a penny looks elliptical I am not mistakenly ascribing elliptical shape to what is in fact round. Sensible appearances *may* lead me to make a mistaken judgment about physical objects, but they *need* not, and, so far as we know, commonly do not. My certainty that the penny looks elliptical exists comfortably alongside of my conviction that it is round. But a mistaken judgment that the penny *is* elliptical would not continue to exist after I knew that the penny was really round. The plain fact is then that 'looking elliptical to me' stands for a peculiar experience, which, whatever the right analysis of it may be, is not just a mistaken judgment about the shape of the penny.³²

Once the point is made, it seems obvious. J.L.Austin, however, in his remarkably influential *Sense and Sensibilia*, insists on taking the

label 'illusion' very seriously. Because 'there is usually no question of anyone being taken in' and 'familiarity, so to speak, takes the edge off illusion'—because, that is, false belief is not generally involved—Austin believes these arguments are flawed from the start.³³ Austin's argument rests on associating *illusion* with mistaken judgement, but, once one realises that 'seems' has a phenomenological as well as a judgemental sense, one can also see that there is a phenomenological sense to 'illusion' that need involve no mistaken beliefs.

The obvious inadequacy of the judgemental approach to 'illusions' was masked by the emphasis those philosophers who were drawn to it put on the language in which perceptual experience is reported, rather than on the phenomenology of the experiences. Support for the theory is drawn from ordinary language. 'Seems', 'looks' and 'appears' idioms definitely have a use in which they express judgements, usually with the implication of uncertainty or denial. We do say such things as 'the economy seems to be improving, but it's not clear yet', or 'he looks well but he is really very sick'. These have the contextual implication that has been called the 'doubt and denial' condition. Those 'ordinary language' philosophers who are doubtful of the meaningfulness of specialised philosophical uses of language have argued that sentences with 'seems' etc. are only used in these 'doubt and denial' contexts and not for reporting 'experience' in general—I use scare quotes, for *experience* is a generic concept of which these philosophers are sceptical.

H.P. Grice argued persuasively that the 'doubt and denial' conditions are not part of the meaning of the 'seems' etc. vocabulary, but only part of what he called 'contextual implicature'; that is, they relate to what someone might reasonably assume about an utterance, not to what is strictly said.³⁴ If, for example, I were asked how many people had been at the match and I replied that there had been more than one hundred, it could reasonably be assumed that there had not been twenty thousand. But if there had been twenty thousand my statement would not have been false, just misleading in the context. Similarly, although to say that something which is plainly and clearly red *appears* red might, in some contexts, mislead by giving an impression of uncertainty, it is not false. Grice's articulation of this commonsense distinction, which only had to be plainly stated to find general acceptance, undercut most 'ordinary language' philosophising. Such philosophy in general rested on the authority of 'what we would normally say', with the assumption

that linguistic practice possessed no deeper rationale which could be investigated. If one accepts a distinction between the semantics proper of an expression and the pragmatics of a particular use—what it informally implies in the context—then the appeal to ordinary use cannot be authoritative about real meaning, for all ordinary use is heavily context-dependent.

If ‘doubt and denial’ is just part of the pragmatics of ‘seems’ and the like, what is its ‘pure’ meaning? Grice pointed out that it is (or can be) to report experience in the ‘phenomenological’ sense—as Broad had remarked forty years before, it ‘stands for a peculiar experience’. Once the phenomenological sense is restored, the argument given above that the weak judgemental theory is inconsistent stands.

The disposal of the weak judgemental theory also tends to undercut another objection to the argument from illusion that was popular amongst ‘ordinary language’ philosophers. This strategy was to deny the existence—or the widespread existence—of illusory phenomena, in the sense required by the argument. Austin, as we have seen, claimed that a stick in water did not look like a bent stick, but like a straight stick *in water*. Similarly, a penny from an angle does not look elliptical, it looks like a round object viewed from an angle. This strategy can be generalised to many, if not all, cases of illusion. If one were taking ‘looks’ in the judgemental sense, then, given that one has no serious tendency to be deceived, then the stick does not look bent and the penny does not look elliptical, and this response to the argument would be powerful. Once, however, one accepts that there is a phenomenological sense to ‘looks’ which is not analysable in terms of a tendency to make judgements, then the situation is more complicated.

I have a fairly clear memory of being perplexed on first noticing what seemed to happen to my arms when I put my hands in the bowl to wash them. I wondered whether it was the upward pressure of the water that bent them, and, if so, why it didn’t hurt. The appearance of bentness is very plain and very convincing. By contrast, the case of the elliptical penny is less clear. Most empiricists have been inclined to treat depth as a matter of judgement and not a bed-rock phenomenon. If visual experience is really two-dimensional, then a partially rotated penny must look elliptical. It is difficult, however, to deny that depth is, to some extent, at least, phenomenally real. It seems to me doubtful

whether a nearby small round object at an angle looks anything other than round.

This perplexity suggests that we need a standard for whether a phenomenon is illusory, in the sense required by the argument from illusion. I suggest the following. An appearance is illusory if the naivest conceptualisation of it would be erroneous: 'naivest conceptualisation' means that conceptualisation one would naturally make, if one had the experience just as it is, but without other experiences (which would mean mainly previous experiences) which would influence one's conceptualisation of it, and without the opportunity for deep thought about what is probable. There is, no doubt, a degree of indeterminacy here, for there will be previous experiences necessary before one could have the experience 'just as it is', and there may be no clear fact—at least, no clearly ascertainable fact—about what exactly these are. But this uncertainty is at too refined a level to affect the point. It is clear, for example, that one would take a colour appearance at face value if other experiences did not lead one to discount it: that someone who had never seen objects more than ten feet away would, on going outside, think that distant objects were smaller: a short-sighted person who had never seen anything close up would think objects fuzzy: and things half in water would be judged bent. In sum, with the exception of some instances of perspective (perhaps, sadly, including the hard-worked rotated penny), the standard cases of illusion pass the test.

3 The difficulty of finding an alternative to the Phenomenal Principle

I remarked earlier that Price's intuitive certainty that when he seems to see something red there really is something red of which he is aware has been given short shrift by more recent philosophers: it has been treated as little better than a howler. Part of my purpose in this discussion has been to suggest that the appeal to intuition in Price, Broad and Moore was in no sense a naive oversight, for the intuition in question is a powerful one to which it is difficult to see a plausible and coherent alternative. In pursuit of these polemical and apologetic purposes, I want to consider a recent account of appearing which is closely related to those here considered. Despite its triumphalist Wittgensteinian tone, Peter Hacker's *Appearance and Reality* constitutes a more rigorous exposition of the approach to perception

associated with Wittgenstein, Ryle, Austin and ‘ordinary language’ philosophers, than is found elsewhere, especially in the originals. The book is principally an attack on the view that secondary qualities are in any way mind-dependent—a theory which Hacker rightly identifies as historically the most influential of the arguments for sense-data. Nevertheless, he has to treat of appearance in general, for the issue of secondary qualities cannot be isolated from the other difficulties for naive realism, especially given that some of these other difficulties are, history apart, more threatening. Hacker is strikingly less convincing when talking about appearance in a general way than when concentrating on the argument from secondary qualities. I say this not primarily as an *ad hominem* remark, but because it exposes the difficulty of saying anything plausible about sensory experience whilst denying the Phenomenal Principle.

Hacker’s treatment of appearance belongs with the theories discussed here, rather than with those that put weight on the idea that sense-contents are *intentional objects*, because he thinks that talk of ‘intentional objects’ means no more than is expressed by ordinary idioms of the ‘seems’ category, and that the latter are primary.

An intentional object or content of perception is not what a person perceives in the sense in which a material object of perception is. The latter can be specified by ostension, the former by a characterisation of *how something perceptually strikes a person*. Hence in cases where a perceptual verb is thus used it can always be paraphrased by ‘It looked...’, ‘It appeared...’, ‘It seemed...’, ‘It felt as if...’, etc.³⁵

The phrase ‘how something perceptually strikes a person’ is later made to do important theoretical work; it is the fundamental characterisation of experience:

‘It looks (smells, tastes, etc.) thus-and-so to me’ expresses how something perceptually strikes one.³⁶

This idiom is not immediately perspicuous. At first sight, it might most naturally seem to suggest the impact of a phenomenal given, but this is not, of course, how Hacker glosses it. He immediately explains it:

Other things being equal, this is what I would take myself to be observing.³⁷

Suppose one is looking at something which is in fact white, but that one has this kind of *prima facie* inclination to take oneself to be observing something red. How is one to understand such an inclination? 'Taking oneself is something of a metaphor, meaning something like 'being inclined to believe about oneself. Is this, therefore, just some kind of belief state—that state of being inclined to believe, *ceteris paribus*, that one is observing something red? For if it is just a state of this kind, as the experience is indistinguishable, in the relevant respect, from veridically seeing red, then veridical seeing must be no more than a kind of belief state, and Hacker is obliged to adopt a reductive theory of perception. If, on the other hand, there is more to it, then what more could this be but the inclusion of some kind of phenomenal object—intentional or otherwise—in the account? But the purpose of Hacker's account was to avoid such objects.

Perhaps this argument is too quick. Hacker's analysis is in terms of taking oneself—that is, having a *prima facie* tendency to believe oneself—to be *observing* such and such, and not simply taking there to *be*—that is, having a *prima facie* tendency to believe there to *be*—such and such, which is what would be required for a reduction of perception to belief. That one can be in a state of believing that one is observing has no implication that observation is merely belief.

This response is inadequate. The suggestion is that observing, for example, red is more than believing, or acquiring some kind of tendency to believe that something red is before one; but that there is another state—namely a taking oneself to be observing—which is only some kind of belief or inclination-to-belief state, which is, or can be, subjectively indistinguishable from it. In that case, that latter sort of belief state is all that is required to be adequate to the phenomenology of perception proper.

4 Generalising the argument

None of the traditional ways of disarming the illusory phenomena on which the argument rests seems to work. The argument, however, is not complete. In my original formulation of the argument I said that *some/much/most/all* cases of perception involve illusion. If objects do not really possess secondary qualities, or are in some other way thought to have been shown by science to be quite different from the way they appear, illusion will be ubiquitous, and the argument from illusion, if

sound, will show directly that we are always aware of sense-data. But if one bases the argument on cases of illusion which contrast with real cases of things looking as they really are, then one will require further argument to show that we are aware of sense-data in the veridical as well as the illusory cases. Broad briefly presents such an argument:

No doubt it would be possible in theory to admit [illusions require sense-data], and yet to maintain that in the one case of direct vision through a homogeneous medium one really is (as one appears to oneself to be in *all* cases) prehending a part of the coloured surface of a remote foreign body. But, in view of the continuity between the most normal and the most abnormal cases of seeing, such a doctrine would be utterly implausible and could be defended only by the most desperate special pleading.³⁸

The ‘continuity’ in question is obvious. There is no absolute distinction between a state of tiredness in which things look slightly less clear and a less tired state, or between accurate vision and very slight short-sightedness; nor, probably, is there such a thing as absolutely accurate perception of colour, rather a slight variation between persons. It is, therefore, very implausible to say that some of these cases involve direct apprehension of an external object and in the others of a sense-datum. So the argument generalises easily.

5 Conclusion

We started the chapter with an informal statement of the argument from illusion. Having considered its various cruxes, we are now in a position to see it laid out more formally. Putting aside the thought that all perception may be illusory (the *argument from secondary qualities* and the *argument from science* will bring us back to it) and taking the uncontentious claim as our first premise, we can state the argument as follows:

- 1 In some cases of perception, physical objects appear other than they actually are—that is, they appear to possess sensible qualities that they do not actually possess.
- 2 Whenever something appears to a subject to possess a sensible quality, there is something of which the subject is aware which does possess that quality.

Therefore

- 3 In some cases of perception there is something of which the subject is aware which possesses sensible qualities which the physical object the subject is purportedly perceiving does not possess.
- 4 If *a* possesses a sensible quality that *b* lacks, then *a* is not identical to *b*.

Therefore

- 5 In some cases of perception that of which the subject is aware is something other than the physical object the subject is purportedly perceiving.
- 6 There is such continuity between those cases in which objects appear other than they actually are and cases of veridical perception that the same analysis of perception must apply to both.

Therefore

- 7 In all cases of perception that of which the subject is aware is other than the physical object the subject is purportedly perceiving.

The assumptions in this argument are (1), (2), (4) and (6):(4), being an application of Leibniz's Law, is not controversial; (6) is hardly controversial, and I have defended it as far as is necessary; (1) has come under rather half-hearted attack from some 'ordinary language' philosophers, but even they did not try to deny that there are *some* cases of things that appear other than they are. It is also challenged by the multiple location theory, according to which objects really do possess all the properties they seem to possess. I have discussed both these approaches. The weight of argument falls on (2)—the Phenomenal Principle. One principal objective of this chapter has been to show the difficulty of getting round (2). I have suggested that gesturing in the direction of the intentionality of 'appears' idioms achieves little or nothing if one is trying to analyse the phenomena. Deeper discussion must await Chapter VII, but I hope to have shown that the casual scorn with which it has become usual to treat the Phenomenal Principle casts more doubt on the judgement of the critics than it does on that of the greater philosophers who accepted the principle.

CHAPTER III

Further Arguments against Naive Realism

1 The argument from secondary qualities

Initial statement and historical examples

We have seen that the argument that we tend to think of as the principal argument for sense-data—the argument from illusion—figured hardly at all in the writings of the early modern philosophers. This, I suggested, was because it did not spring in any special or direct way from the atomism of seventeenth-century science. Among the arguments that were rooted in science perhaps the most important rested on the distinction between primary and secondary qualities. The argument can be sketched informally as follows:

Science has shown that physical objects do not possess secondary qualities intrinsically. As they are clearly possessed by that of which we are aware in perception, that of which we are aware in perception is not the physical object itself. The only plausible way to understand the relation between physical objects and secondary qualities is to think of the objects as possessing dispositions to produce the qualities in us as properties of our sense-data.

The idea that objects do not possess secondary qualities intrinsically has a long association with atomism. The ancient atomist Lucretius described secondary qualities as existing only ‘by convention’.¹ Admittedly, there is no agreement about exactly what this means, but it certainly means that they are not a full part of the nature of physical things.

Examples of a similar but clearer doctrine can be found throughout the writings of seventeenth-century philosophers and scientists. Galileo, for example, is quite explicit on the matter:

Now, whenever I conceive of any material or corporeal substance, I am necessarily constrained to conceive of that substance as bounded and as possessing this or that shape, as large or small in relation to some other body, as being one, many or few...But I do not at all feel myself compelled to conceive of bodies as necessarily with such further conditions as being red or white, bitter or sweet, having sound or being mute, or possessing a pleasant or unpleasant fragrance...I think, therefore, that these tastes, odours, colours, etc., so far as their objective existence is concerned, are nothing but mere names for something which resides exclusively in our sensitive body, so that if the perceiving creatures were removed, all of these qualities would be annihilated and abolished from existence. But just because we have given special names to these qualities different from the names we have given to the primary and real properties, we are tempted into believing that the former really and truly exist as well as the latter.²

Descartes expresses the same sentiment:

[N]othing whatever belongs to the concept of body except the fact that it is something which has length, breadth and depth and is capable of various shapes and motions...But colours, smells, tastes and so on, are, I observed, merely certain sensations which exist in my thought, and are as different from bodies as pain is different from the shape and motion of the weapon which produces it.³

Locke's statement is probably the most famous.

[T]he ideas of primary qualities of bodies are resemblances of them, and their patterns really do exist in the bodies themselves, but the ideas produced in us by these secondary qualities have no resemblance of them at all.⁴

The particular bulk, number, figure and motion of the parts of fire or snow are really in them...But light, heat, whiteness, or coldness are no more really in them than sickness or pain in the manna. Take away the sensation of them...and all colours, tastes, odours and sounds, *as they are such particular ideas*, vanish and cease, and are reduced to their causes, i.e. bulk, figure, and motion of parts.⁵

Possible responses

The idea that objects do not possess secondary qualities intrinsically, rests fundamentally on the thought that such qualities do not figure in the basic scientific—that is, the physicist's—account of the world. This is because scientific explanations of events always operate in terms of primary qualities: as secondary qualities are causally idle no purpose is served by attributing them to objects intrinsically.

I shall consider five responses to this argument, but the fifth must await discussion of the 'argument from science'. According to two of these responses, it is wrong to say that secondary qualities are inert: as they are, contrariwise, active, there is every reason to accept common sense and attribute them to objects intrinsically. One of these responses is reductive and the other is not. The reductive response says that secondary qualities are active because they are in fact identical with certain primary quality structures. This is the sort of view that says, for example, that phenomenal colour just *is* light-waves at the surface of objects, and that phenomenal sound is waves in the air. The non-reductive theory rejects the identification of the quality with the underlying primary qualities; the phenomenal colour is not the same as the light-waves—though no doubt they are closely connected—but the irreducibly phenomenal does possess causal powers.

The third response to the argument is to accept that secondary qualities are inert, but to deny that it follows from this either that physical objects do not possess them intrinsically or that we do not perceive those qualities directly: naive realism is untouched by the physical inertia of secondary qualities. The fourth response accepts the conclusion of the argument—that secondary qualities in objects are only dispositions to cause us to have experiences of certain kinds—but denies that this is incompatible with naive realism. The fifth option is to import talk of 'levels' and say that secondary qualities are part of the 'manifest image' of the world, which is just as real as the 'scientific image' and not in competition with it, these being different 'levels' of reality or of description. Because of its connection with the 'argument from science', this objection will be discussed when I discuss that argument.

Because it is nearest to common belief I shall consider first the second response, which was that secondary qualities are both irreducible and causally active.

That secondary qualities are irreducible and causally active

Peter Hacker's book *Appearance and Reality* is primarily an extended attack on the attempt by philosophers, from the seventeenth century to the present day, to use secondary qualities as the basis for the attack on common-sense realism. He rejects reductionism and yet holds that secondary qualities have causal and explanatory power.

We correctly explain why my hut is cool in the summer while yours is hot by reference to the fact that yours is black and mine is white. For black objects absorb light and warm up in the sun to a far greater degree than white ones. That *this* fact is *further* explained in terms of the molecular or sub-atomic structure of the white and black surfaces does not in any way show that the fact that a surface is black does not explain why it warms up... Indeed, if it did, then by parity of reasoning, the fact that objects are solid, liquid or gaseous would also, *mirabile dictu*, be explanatorily idle, since these properties too are explained in terms of molecular structure.⁶

Hacker's argument is to take as a premise our ordinary belief that colours interact with light, and then reduce to absurdity the Lockean interpretation of this—which is that it is really microscopic primary qualities of surfaces that interact with light, not secondary qualities such as colour—by saying that parity of reasoning would show that macroscopic primary qualities are equally inert.

The problem with this argument is that there is no parity in the reasoning. The micro-structures that make an object solid, liquid or gaseous are *a priori* sufficient for the macroscopic properties: that is, it can be deduced from the properties of the micro-structures and the laws governing their behaviour on the microscopic level that the macroscopic conduct of objects so constructed will be that of a solid, a liquid or a gas. There is a simple impossibility in the microstructures doing what their sciences say they do and the macroscopic object turning out to have the wrong property. If the molecules bind in a certain way the object just cannot, for example, behave as a liquid. This is not an empirical truth, for if the molecules are binding tightly, that *means* that they don't move easily relative to each other, which *entails* that the object is not flowing. We can say, therefore, that the macroscopic causal powers of, e.g., solidity, are *constituted* by whatever structure realises it in a given case and, hence, are not threatened with redundancy by it.

The situation for primary qualities is, therefore, parallel to that between colours and structures on a *reductive* theory of colour. On Hacker's non-reductive theory there is, *ex hypothesi*, no *a priori* connection between structure and colour and so no simple impossibility of the occurrence of that structure with a different—or no—phenomenal colour. The structure does not, in the same way, simply constitute the colour and so the efficacy of the structure does not constitute the efficacy of the colour. As the structure is itself causally sufficient for all effects, and the colour is something further than the structure, in the sense of not being entailed by it, there appears to be no causal role for the colour to occupy.

***That secondary qualities are to be understood reductively
and are causally active***

According to D.M. Armstrong and others, when we perceive secondary qualities we are in fact perceiving primary qualities in a confused, indistinct or blurred way.⁷ What we recognise when we see something as red is that something is going on at that object which is like what goes on at the surface of post-boxes, fire-engines, tomatoes, etc., but we cannot tell just by seeing what that thing 'going on' is: in seeing things as red, we make *topic-neutral* judgements that certain things resemble each other.

This theory is usually—if not always—associated with a theory which reduces perception in general to the acquisition of belief, or information; that is, a theory which denies that there is in experience anything which is not capturable in terms of propositional content. This association is not difficult to understand. There is no difficulty in seeing how the content that 'something is going on like what goes on on the surface of pillar-boxes, etc., but I cannot tell what it is intrinsically' can be captured propositionally—I have just done it—but it is difficult to see what this would mean if one tried to understand the phenomena non-reductively. In either sense-datum or naive realist terms, the theory would come to the claim that when one experiences, e.g., a colour, one is aware of an expanse which is characterised by no positive quality but which is in some inscrutable respect like various other expanses. Nevertheless, a proponent of a non-reductive version of this topic-neutral approach to secondary qualities might reply that they were merely providing a 'resemblance theory' analysis of secondary-quality universals, whilst allowing a

properly realist account for primary-quality universals.⁸ I do not find this suggestion easy to judge, but am not inclined to think it plausible. Once one has conceded that a realist account is the correct one for genuine properties (as the reductionist about secondary qualities must, for if he or she applies a resemblance analysis to primary qualities as well, then they are no more real than secondary ones) then one will probably feel no more inclination to think that the resemblance analysis does justice to colour than to think it does to shape. If properties in general are not to be analysed in terms of similarity-dissimilarity relations, how credible is it to conceive of a visible square of colour as a real square patch occupied simply by a certain kind of likeness-to-this and unlikeness-from-that? This seems to make sense only in the context of judgements or beliefs—that is, in the context of a reductive strategy about the experience as such.

I shall be discussing reductive theories of perception in Chapter V, and, in so far as this treatment of secondary qualities is dependent on those theories, the final discussion must await that chapter. There are, however, powerful empirical objections to all available reductive accounts. Physical theory suggests only two plausible candidates for being identical to colour, namely light wavelength and spectral reflectancy of surfaces. C.L.Hardin shows convincingly that neither of these features correlates sufficiently with our standard colour perceptions.⁹ The core of the problem is that colour perception depends on only three types of receptor in the eye ('cones'). Each of these cones responds to a particular range of wavelengths and is either firing or not—that is, it does not respond variably to the strength of the stimulus. An enormous amount of information is, therefore, lost at the eye, which has only very limited resources to map external physical features. Consequently, quite different physical features can 'play the same tune' on the cones of the eye and, hence, present themselves as coloured in the same way. There is, therefore, no useful unity in the set of wavelengths that look to be the same colour, even in quite standard situations. Surface reflectance fares somewhat better in this respect, but entirely leaves out of account those colours that do not operate by reflection. This includes not merely such exotic things as holograms, but also television pictures. Even for colours that are reflected, what colours they appear as depends on contextual features, such as what other colours are being presented at the same time. It is important to notice that all these variations occur in what we would think of as normal situations, rather than tricks or illusions. So the only physical

processes with which normal colour appearances are reliably correlated are internal to the perceiver's perceptual system. Colour perception is not, therefore, a blurred or topic-neutral perception of specific external physical conditions.

Secondary qualities are inert but intrinsic and perceived

With the reservation that some of the argument requires one to skip forward, I have defended the view that secondary qualities are inert. Accepting this, an opponent of the argument might claim that objects could possess secondary qualities even if they were causally idle. There might be no scientific point in attributing them to objects, but there might be metaphysical point, for, in abstraction from secondary qualities, it might be argued, we can make no sense of physical objects.

This possibility forces a major revision of the argument but does not affect its essential drift. The problem is that if intrinsic secondary qualities are causally idle they are, *ex hypothesi*, idle in the production of our experiences: and it seems plausible to claim that something which has no role in producing an experience cannot be what one is aware of in having that experience. In other words, if the quality is idle in the production of the experience, the experience could have taken place in its absence and been exactly similar; so the quality is not what one is aware of in the experience. The argument can be recast:

Science has shown that physical objects possess no causally active and intrinsic secondary qualities. Something that is not causally active in producing an experience cannot be what we are aware of in that experience, so no secondary qualities intrinsic to physical objects are among what we are aware of in perception. As we clearly are aware of secondary qualities in perception and they must belong intrinsically to something, then what we are aware of in perception is something other than a physical object. The only plausible way to think of the relation between external physical objects and the secondary qualities we perceive is to think of the object as being disposed to produce sense-data which instantiate those qualities.

This argument is not as strong as it seems at first sight. The new premise is that something which is not causally active in producing

experience cannot be part of the content of that experience. This cannot, for a sense-datum theorist, be a universal truth, for sense-data themselves, it would seem, partially *constitute* rather than *cause* experiences. The thought is probably that something *external* could not be perceived if it were causally idle, for we would not perceive it unless it made some sort of impact on us. I think this idea is confused. Some theories of perception—for example, representative theories and at least some reductionist physicalist ones—require a causal criterion for what is to count, from amongst the features of the external world, as the features perceived. Plainly, within such a framework, no causally redundant feature could be perceived. But the situation is less clear in the case of a non-reductive direct realism, because, in naive realism, the external object plays the same role as the sense-datum does for the sense-datum theorist, by being not so much the cause of the experience as its principal constituent. Consequently, under naive realism it is not very clear what is meant to be the exact nature of the relation between the causal process and the experience.

There are two possible theories that one can hold about the connection between the process and the experience. According to one, the brain state in which the causal process ends up is sufficient to produce the perceiving together with its subjective content. This—dubbed the *generative theory*—is essentially the rationale of the causal argument against naive realism. On the other theory the causal process produces the *act* of perception, but the content is constituted by the features of the external world on which that act is directed. This has been called the *selective theory*, because the causal process enables us to pick out the content, it does not generate it.¹⁰ To avoid the causal argument for sense-data the selective theory is the one that a naive realist must adopt. Someone trying to argue for sense-data *on the basis of the nature of secondary qualities*, cannot let the argument rest on an appeal to the generative account of the physical process, for then the argument would already presuppose something which could be independently used to establish the conclusion—the appeal to the nature of secondary qualities would cease to be important. Only if an independently justifiable account of secondary qualities entailed a generative understanding of the causal process could that understanding be appealed to without, in effect, begging the question. If one sticks to the selective account of the causal process, then there is

no reason why something which is causally inert should not be the object of perception.

The point is that even though, for example, the colour in the object may not causally affect what light-waves are reflected from it, the naive realist could still claim that it is that very colour-instance that is apprehended in the act of perception to which those light-waves finally give rise. The causal process triggers our visual consciousness which falls directly on the colour 'out there': the colour need only be at the right place, defined by the causal process itself; it is irrelevant whether it plays a causal role.

This naturally raises the problem of what would have happened if those light-waves had been reflected by an object with a different intrinsic colour, or none at all, which should be possible if the colour played no role in the reflection. A naive realist cannot say that the experience would have remained the same, for there would now be nothing for the consciousness to light upon and make into its content. There are two things that they could say.

The naive realist might argue that the subject would perceive whatever colour was in fact correlated with the light-waves at the object: so if there were none they would perceive nothing. The rationale for this is that the function of the causal process is simply to prompt a perceptual act directed at the object from which the process emanated, and not to determine what consciousness finds there. There are grave difficulties with such a theory. For example, why should consciousness land on the object that reflected the light and not at some other point in an essentially continuous causal process? But this difficulty faces the naive realist anyway, for it necessarily accompanies the view that the mechanism causes act but not content of perception. There is also the problem of how, on this theory, misperceptions should take place, but that is the argument from illusion and owes nothing special to the role of secondary qualities.

The second strategy is to claim that there is a set of brute natural laws that correlate certain surfaces (or structures, if one wishes to think of colours as pervasive throughout) with certain intrinsic colour qualities. This is neither more nor less arbitrary than the laws an epiphenomenalist thinks holds between brain and mental states—which is not surprising as the theory we are discussing is a sort of epiphenomenalist theory of secondary qualities (only 'sort of' because the qualities, by being contents of consciousness, do affect minds, though they have no direct effect on matter). It would then be part of

the design of the world that the right qualities were always available to be perceived. This theory faces the same problems concerning the causal process and illusion as the first, but the reply is the same; the naive realist faces these problems anyway and secondary qualities play no special role in them. (Though this reply is conceptually legitimate, it does, of course, run foul of the fact, cited above, that there are no external physical conditions that are reliably correlated with the appearance of a particular colour.)

The upshot of the argument so far is that the causal idleness of secondary qualities does not entail that they are not intrinsic to physical objects, and that if they were idle but intrinsic this would create no problems that the naive realist would not have to face on quite independent grounds. It does not follow from this that there is no argument based on secondary qualities against naive realism. Nothing said so far undermines the conditional that *if* objects do not possess secondary qualities *then* naive realism is false. All we have is the reservation that causal idleness does not entail that they are not intrinsic. It is open to us either to provide some further reason for thinking that they are not intrinsic, or to argue that, although idleness does not entail non-intrinsicity, it makes it unreasonable to believe in intrinsicity.

There are three reasons for preferring the theory that secondary qualities are intrinsic, *ceteris paribus*.

- 1 It accords with common belief.
- 2 It is necessary for naive realism, so, in the context of an attempt to refute naive realism convincingly, the benefit of the doubt should go to the intrinsicist.
- 3 It is doubtful whether the concept of matter can make sense unless one attributes secondary qualities, or qualities analogous to secondary qualities, to matter: a matter which consists of primary qualities only is too formal and bare.

Against these considerations can be urged the peculiar dependence of secondary qualities on particular kinds of experience. Shape, for example, can be directly experienced (that is, its presence is not merely inferred) by either touch or sight, and these two forms of experience are subjectively very different. But colour or sound can be directly experienced only through experiences which are subjectively like sight and hearing. It is a mark of objectivity that something can be approached in a variety of ways, because, if something is objective, its

existence is not conditioned by how we experience it. So secondary qualities are not objective, but depend on the nature of the senses and are, hence, not intrinsic to external objects. This argument is reinforced by the consideration that, as secondary qualities are tailored to match senses, a proliferation of senses would lead to a proliferation of secondary qualities. So if we imagine varieties of creatures with different sense organs—perhaps we could imagine developing a new range of such creatures—there would have to be a secondary quality for each sense. This kind of dependence on senses makes intrinsicity seem implausible and contrived.

This argument has much less force than it seems to have. It contains two relatively distinct points. The first is the sense-relativity of secondary qualities and the second is the possibility of the proliferation of such qualities virtually *ad lib*. The sense-relativity can be seen, however, not as a sign of the dependence of the quality on the sense, but as a reflection of what it is for something to be a quality of the relevant kind. Shape is, after all, a rather formal kind of quality, but colour is more gutsy and less abstract. It is not sight that determines the nature of colour, but colour that determines the nature of sight: it is because of what colour is like that an experience would only count as a direct experience of colour if it were subjectively like seeing.

Nor is it obvious that secondary qualities can be proliferated *ad lib*. Not just any neural structure gives rise to sensory-type experiences, but only (it seems) ones that normally become active in response to certain kinds of external stimuli. If different kinds of creatures are responsive to different kinds or ranges of radiation from those that we can sense, then they may experience different secondary qualities, but why should there not be other qualities associated with stimuli that we do not detect? There is nothing arbitrary about this. The real thought is that the same stimuli could give rise to different experiences if they had been associated with different neural structures. If, for example, the human eye had been connected to those centres which are in fact associated with hearing, light would have given us sound experiences. Other creatures might have different neural arrangements associated with light so that their experiences might be goodness-knows-what. Which arrangement evolution might throw up is sufficiently arbitrary for there to be no plausibility in describing some of these as right and others wrong, so it cannot be right to attribute any of these qualities to the objects themselves.

Why should we assume that, had the eye been originally associated with the auditory centres, we would have heard light? This can only be because we assume that it is a natural and inalienable property of a certain kind of neural structure that it produces a certain kind of experience. The sense-datum theorist, who is a kind of dualist, faces a paradox here. Locke pointed out that it is entirely arbitrary that a certain brain structure should give rise to one sort of experience rather than another—or than none. Modern research on the brain confirms the fact that, cell by cell, it is very homogeneous. There is no reason, therefore, why, on the strength of its intrinsic nature alone, a particular area of the brain should have been devoted to a particular kind of stimulus. This fits in very well with the naive realist's commitment to a *selective* not a *generative* account of the role of the causal processes in perception, for this means that the process does not generate a content but puts one in touch with the stimulus that is out there. Any neural state which was able to respond differentially to the elements in a stimulus would be able to perform that role. I can see no reason, therefore, why a naive realist should not say that our auditory centres could have been dedicated to vision and the perception of colour, if that had been the stimulus to which they had originally been dedicated. Of course, the rewiring of our brains in mid-life, when an association between those processes and colour had already been set up, would doubtless give rise to some very inappropriate experiences, but that would be the result of the change, not of their intrinsic inappropriateness to be the vehicles of vision.

The theory that secondary qualities are not intrinsic to physical objects is part and parcel of a certain kind of scientific picture of the world. That picture has a powerful appeal and is incompatible with naive realism. One of the respects in which it is incompatible with naive realism is that it denies the intrinsicity of secondary qualities. But any argument for treating secondary qualities in this way is going to rely on a version of the causal theory of perception which already presupposes that one has disposed of naive realism, and adopted a representative theory. Once having reached this point, it becomes attractive to deny the intrinsicity of secondary qualities, which reinforces the conclusion about naive realism that one has already reached. But there is not here an independent argument against naive realism.

Dispositional direct realism

The fourth strategy open to the opponent of the ‘argument from secondary qualities’ is to accept its conclusion but deny that this constitutes an abandonment of naïve realism. The essence of this position is to reconcile two apparently inconsistent propositions. On the one hand, it accepts that colour in objects in the absence of the perceiver is no more than a structurally grounded disposition to produce experiences of a certain kind in us—to make us see the object in a certain way, namely as coloured. On the other hand, it affirms that the colour that we see the object as possessing really is ‘out there’ on the surface of the object when we are seeing it. Colour is just a disposition of the object to make us see it as coloured—a disposition that rests on how we are constructed as well as on the nature of the object and the medium—but then we really do see *it* as coloured, in a literal sense; that is, it is not a matter of having a coloured sense-datum which we *see as* or *project onto* the object.¹¹

The natural response to this suggestion is puzzlement as to how, if colour is, or is a function of, an effect produced in us, it can exist in external space, except by some sort of intentional projection. The answer to this is, I think, meant to be that its existence in public space is a matter of its being the content of an intentional object, but it is not a consequence of projection.

The projectivist view of some property involves maintaining that the property is *actually* present in some psychological state, but then is spontaneously experienced *as if* it were in the external world—it is projected onto external objects. On the traditional, Lockean interpretation, this is what happens to colours. This account will only work if it makes sense to think of the property in question as *really* qualifying some mental state. In the case of colours, this means colours characterising some private visual field. This, too, is the traditional account. If one believes either that there cannot be anything private and spatial, or that there cannot be anything private at all, then this projectivist account will not be available. Then, if one sticks by the dispositional account, colour will exist only when seen as a property of something external: one could still say that it is projected, but that it exists only in its projected form; that which is projected is, in its pre-projected form, something less than colour as we know it in experience. It might be compared with Kant’s given, which, prior to being brought under the categories, is only proto-experiential.

Any theory which says that an object's being red is just a dispositional property to make us see it as red, with phenomenal red existing only in our seeing, has a problem about the instantiation of phenomenal red. On the one hand, phenomenal red—red-as-we-experience-it—is not instantiated in external things, for, *ex hypothesi*, they are red only in the dispositional sense. On the other hand, it cannot be instantiated in experience, for that would mean that sensible qualities are instantiated in experience, and that is the sense-datum theory. It follows that colours are only intentionally inexistent—that is, that there are no instances of colours, either in physical objects or sense-data. This theory is very bizarre, but it can be finally refuted only by refuting the intentionalist account of perception. Someone might try reconciling dispositional direct realism and pure intentionalism by holding that, just as hobbits were never intended not to be fictional, so colours are only supposed to be intentionally inexistent; it is just a mistake to think that colours are the sorts of things that are supposed to have instances. This is because they are only supposed to be features of how things appear, not how anything is. Once again, the pressure will be on the adequacy of an account of intentional objects.

The best thing that a dispositional direct realist can do if they are to avoid pure intentionalism is to claim that, in the appropriate cases, red is instantiated, because causing one to see it that way is what it is to be red. For an object to have the disposition to cause one to see it as red is what it is for it to be red and, therefore, the red that it appears to possess really is instantiated in it, just in that sense. My suspicion here is that this gives a sense to *being red*, namely, *causing one to see it looking like THIS*, but that this gives no sense to the instantiation of the property designated by *THIS*, which is the phenomenal property we wanted to be able to think of as instantiated.

Even if this were not an insuperable problem, I think that there are serious difficulties with conceiving of public phenomena as being dependent on the way subjects are affected. If I am to conceive of the red patch of which I am aware as the product of a dispositional influence on me, how can I think of it as the same red patch as that of which you are aware? And if the patch gets smaller as I retreat from an object, how can it be identical with the surface of the object that remains the same size? A non-dispositional direct realist would say that the patch does not get smaller—it only seems to, but this reply lacks conviction from the dispositionalist, who sees the phenomena as

relative to how he or she is affected, and certainly this varies in magnitude with distance.

I think that these things suggest that dispositional direct realism is not a factual thesis, but a grammatical one, that seeks to relate the discourse of naïve realism to the ontology of scientific realism. The theory as I have already expressed it contains three points: (1) it is a feature of our discourse that we attribute colours to external objects; (2) it makes no sense to attribute them to a subjective or private state; (3) colours are observer-relative, that is, they are, in some sense, dispositions in objects to affect observers. To deal with the problems just raised one would add (4) that we *say* that it is the same patch that we all see and that the patch is the size of the relevant surface; and if it looks different shades or sizes to different people we *say* it *looks* different but *is* some standard shade or size. We also *say* that that is what it is for an object to be red and, therefore, for red to be instantiated. Whereas a traditional empiricist would seek to examine language critically in the light of the realities of phenomenology and science, some moderns take the ‘logic’ of the language as being—or, at least, enshrining—all there is to reality. This is not, in the present context, a dogmatic affirmation of ordinary language philosophy. The attack on the private forces one into a direct realist construal of the phenomena, and prevents one from assigning any ontological standing to non-veridical phenomena. In practice, this means that there is no more basic reality to experience other than that expressed in ‘what we would say’, for any more basic reality could only consist in the realm which was ruled out by ruling out the private. If the attack on privacy can be shown to be groundless, then we shall lose the only reason for taking language as the ultimate authority on experience, and be able to take the phenomena as authoritative. The idea that experience is somehow a grammatical phenomenon—a creature of the ‘logic’ of our language—seems to me an absurdity; but it is worth noticing that it is an absurdity that goes hand in hand with the denial of the private. Both these issues will be discussed later.¹²

Conclusion about the argument from secondary qualities

The view that secondary qualities are causally active because they are identical with primary qualities in objects seems to be empirically false, and also rests on a reductionist account of perception which will be discussed—and refuted—later. Non-reductionist activism goes against

science and Hacker's attempt to save it rests on a mistake. Dispositional direct realism is a subtle theory but is sufficiently in tension with the phenomena to require the support of Wittgenstein's attack on privacy: the failure of this attack will be shown in the next chapter. The view that secondary qualities are intrinsic but inert, however, is not refutable without independent arguments for sense-data. Though not strictly refutable, it is, nevertheless, a very odd theory, requiring the existence of brute laws that correlate phenomenal properties with the mechanisms that cause us to perceive them. It seems to me that the argument, though not conclusive, is fairly persuasive.

2 Argument from science stated

The argument from science can be seen as a special case of the argument from illusion or as an extension of the argument from secondary qualities. The foundation of the argument from illusion is the premise that, in ordinary perception, things tend to appear different from the way they actually are. In its ordinary version, the thought is that objects look the way they are only under the best circumstances of observation. The argument from secondary qualities bolsters this up by arguing that a whole category of sensible qualities is such that perceiving objects as they really are in these respects is not, in a naive realist way, even a possibility, for objects do not possess secondary qualities and one can never perceive an object as lacking all secondary qualities. Science, however, seems to go further and to suggest that nothing looks the way it really is even in respect of primary qualities. Microscopic observation shows, in particular, that no edge is really smooth and that so-called solid bodies are largely empty space. Modern—as opposed to Newtonian—science seems to go further and to show that apart, possibly, from spatial location, the properties of bodies are of an utterly different kind from any which we can perceive. The argument can be summarily stated as follows:

Objects appear to possess the usual range of primary and secondary qualities. Science shows, however, that, apart from spatial location, objects possess none of the properties that figure in the 'manifest image' of the world. So virtually all sensible qualities (which make up the 'manifest image' of the world) are

creatures of perception; that is, they exist only in subjective states.

Some classic sources

This argument, in its Newtonian form, ranks with the ‘argument from secondary qualities’ as one of those important in early modern philosophy. It is, like that other argument, intimately associated with the atomist picture of the world, reinforced by the invention of the microscope. Locke is mainly concerned with how this affects the secondary qualities that we perceive, though this involves the discovery of new shapes.

Had we senses acute enough to discern the minute particles of bodies, and the real constitution on which their sensible qualities depend, I doubt not but they would produce quite different ideas in us...This microscopes plainly discover to us ...Blood, to the naked eye, appears all red; but by a good microscope, wherein its lesser parts appear, shows only some few globules of red, swimming in a pellucid liquor, and how these red globules would appear, if glasses could be found that could yet magnify them a thousand or ten thousand times more, is uncertain.¹³

Locke’s science still preserves the primary qualities of common experience, but modern science is more radical. The classic statement of the modern view is found in Eddington:

I have settled down to writing these lectures and have drawn up my chairs to my two tables. Two tables! Yes; there are duplicates of every object about me—two tables, two chairs, two pens...One of them has been familiar to me from earliest years. It is a commonplace object of that environment which I call the world...It has extension; it is comparatively permanent; it is coloured; above all it is *substantial*...Table No. 2 is my scientific table. It is a more recent acquaintance and I do not feel so familiar with it...My scientific table is mostly emptiness. Sparsely scattered in that emptiness are numerous electric charges rushing about with great speed; but their combined bulk amounts to less than a billionth of the bulk of the table itself...There is nothing *substantial* about my second table. It is

nearly all empty space—space pervaded it is true by fields of force, but these are assigned to the categories of ‘influences’ not of ‘things’.¹⁴

There are four strategies used against this argument. One is the paradigm case argument; the second is the view that macroscopic and microscopic properties are simply compatible; the third is they are compatible because they exist from different *perspectives*, and the fourth is a variation on the dispositional realism discussed in connection with secondary qualities.

3 The paradigm case argument

This argument is that our concepts of, for example, solidity and well-definedness, are introduced and explained by reference to certain paradigm objects: it therefore makes no sense to say that these objects are not solid or well defined, for they establish what we mean by those words.¹⁵ Nothing further that we discover about them can show that they are not that of which they are paradigm instances, but only what the ‘inner workings’ are of possessing those properties.

The paradigm case argument is correct about solidity in the following way: no scientific advance is going to prove that two ordinary billiard balls really pass through each other when they collide. In so far as solidity is understood as a behavioural or dispositional property, the objects we take to be solid are solid. The idea that matter is not fundamentally solid is not trading on the dispositional notion of solidity at all. It rests on the ‘solar system’ model of the atom. Just as the planets spin round the sun in a vast empty space, so, proportionately, electrons spin round the nucleus. It is the picture of ‘empty space’ that is doing the work here. We naturally associate empty space with ease of penetration and so this idea is naturally associated with not being solid in the dispositional sense. Solidity is being thought of as a quality which *occupies* space continuously throughout the solid object, thus ruling out the possibility of empty space in that region; it is also at the same time the ground for the disposition. There are well-known problems with trying to combine the quality and the disposition in this way. But even putting aside these problems, the ‘empty space’ picture is misleading. It suggests that it is easier for things to pass through atoms than it is, and it rests on a very particulate notion of the elementary particles.

Solidity, then, if understood in a reasonably sophisticated way, is susceptible to the paradigm case argument, because of its dispositional nature. This is not so for non-dispositional properties. Microscopes reveal that smooth surfaces are rough and straight edges crooked. These notions can be explained dispositionally, so that an object is smooth if it will slide in a certain way or if it feels smooth to the touch. But there are also more absolute geometrical notions in play here. A microscope reveals a surface to be rough in just the same sense of that term as that in which a rough macroscopic object is rough. If our eyes had been better or our fingers more sensitive we would never have thought the surface smooth and we would never have thought the edges straight.

4 Simple compatibilism

The simplest way of trying to reconcile them would be to claim that science and perception just happen to latch on to largely different features of the world: objects simply possess both sets. This line fits well with the paradigm case approach to solidity: the sort of macroscopic object we normally call 'solid' can be made of objects that we would not call 'solid' on their own. However, if what I say above about 'straight edged' and 'smooth' is true, then the compatibilist approach will not work there, because the properties will be inconsistent. It is difficult to see how the difference in resolution involved in these cases can be explained except by invoking the relation to the perceiver. This in fact brings us to perspectivalism.

5 Perspectivalism

One response to the difference between the manifest image of the world and its scientific image is to say that they do not clash because they represent different perspectives on the world. An object cannot be *F* and not-*F*, *simpliciter*, but it can be *F* from one perspective and not-*F* from another: in this case, objects can be coloured from the common-sense perspective and not coloured from the scientific perspective. Variants on this approach are taken, for example, by both Dummett and Strawson.¹⁶

There are two versions of the theory, one in which 'perspective' is taken literally and one in which it is metaphorical. Dummett has a literal version of the theory. He distinguishes between descriptions *in*

absolute terms and descriptions *in relative terms*. The latter say how things are from a certain viewpoint and the former, which is the perspective of science, tries to eliminate observer-relative or viewpoint considerations. The viewpoints involved in the relative descriptions are literally spatial viewpoints: they say how an object appears to an observer of a certain sort in a certain location. I considered in Chapter II attempts to solve the problem of illusion by relativising properties to perspectives in this kind of way, and, if what I said there is correct, sensible qualities cannot be understood in this way. Dummett simply assumes that the notion of an object's looking—or being—F from here and not-F from there is unproblematic for all normal sensible qualities. So does Strawson, but, when he applies the idea to the relation between manifest and scientific image, the perspectives have ceased to be literal, and are concerned with 'frameworks' and 'standards' rather than locations:

We can shift our point of view within the general framework of perception, whether aided or unaided by artificial means; and the different sensible-quality ascriptions we make to the same object are not seen as conflicting once their relativity is recognised. Can we not see the adoption of the viewpoint of scientific realism as simply a more radical shift...? ...Of course, the scientific point of view is not, in a sense, a point of *view* at all. It is an intellectual, not a perceptual, standpoint.¹⁷

Strawson, like Dummett, thinks it unproblematic to invoke perspectives to solve the problem of incompatible apparent properties in actual perception. We have seen that this is wrong; but it does not follow necessarily from the fact that complacency about the perceptual case is misplaced that the metaphorical or 'intellectual' version fails. However, the metaphor of intellectual viewpoint is not going to look as if it solves the problem of incompatible properties if literal viewpoint fails to work, for it is supposed to draw its plausibility from that case. Without this support, talk of 'intellectual viewpoints' looks very suspect. For Strawson is claiming that it is acceptable to attribute seemingly incompatible properties to an object, provided that they belong to different intellectual perspectives, *just as it is possible to do so for literal perspectives*. In so far as the problem is not solved for ordinary perspectives, one is simply left with the ascription of incompatible properties. He admits that 'Such a relativistic conception will not please the absolute-minded'.¹⁸ One's

'absolute-mindedness' need go no further than a distaste for contradiction.

6 Dispositional direct realism and the manifest image of the world

In so far as the properties of the manifest image are, according to this reading of science, all of them absent from the world as conceived by science, then they are all in the predicament normally imputed to secondary qualities. One might, therefore, think that if dispositional direct realism works for secondary qualities it might also rescue the apparent primary qualities of objects. Clouds of electrons, even though not actually square or solid, have the disposition to look square and solid, just as they have the disposition to look red or yellow. If red and yellow cannot be properties of private mental states, neither can squareness and solidity, so they could equally be dispositional yet really external properties of objects. The situation is complicated for primary qualities, however, by the fact that there appears to be a positive contradiction between those imputed by science and those in the manifest image. An object cannot have both a perfectly smooth edge and the rough edge revealed by a microscope. Somewhat ironically, perhaps the properties of modern physics are so different from those of Newtonian science, that there is a difference of category rather than simple contradiction. Even so, the same trick cannot be worked as can be tried for secondary qualities. Whereas *being coloured* can be argued to be *being disposed to appear* a certain way, being square etc. has an independent geometrical definition. So one cannot argue that to be square is to look the way certain things look from head on: the question of whether the objects that normally look that way actually instantiate the geometrical property can still be raised, and the answer will be negative. It follows that the primary qualities in the manifest image can be no more than intentional objects, for looking like that is not what it is to be square; the whole of the manifest image of the world will be intentionally nonexistent. This is the same situation as trying to solve the argument from illusion in its most general form by appeal to intentionality, with the extra counter-intuitive thought that the intentional objects seemingly present properties of radically different kinds—as opposed to differing in their instances—from the external world. This can only finally be judged when we consider the intentional approach to experience.

7 The time-lag argument

The time-lag argument has a restricted and a general version: the restricted version can be taken alone or as preparing the ground for the general version. The argument usually starts from the observation that the distance from some stars is so great that, by the time we see them they have ceased to exist. This piece of information and the data about the speed of light are more recent than the classical empiricists: this, therefore, is an argument that did not appear before the nineteenth century. Slightly less shocking than the non-existent but visible stars is the fact that light takes eight minutes to arrive at the earth from the sun, thereby enabling us to see how the sun was eight minutes ago, not how it is as we look at it. Then the fact that all causal processes take time is invoked to generalise from these astronomical phenomena to the fact that we always see things as they were in the however-so-recent past. The restricted version of the argument runs as follows:

- 1 It is an undisputed fact that, in some cases, perception takes place after the physical state perceived has ceased to exist.
- 2 The content of perception must be contemporaneous with the perception of which it is the content.

Therefore

- 3 In some cases, the content of perception exists after the physical state perceived has ceased to exist.
- 4 If one thing exists at a time at which another does not then they are not identical.

Therefore

- 5 In some cases, the content of perception is not identical to the physical state perceived.

The generalised version of the argument is slightly different.

- 6 All perception involves a temporally extended causal process which begins at the object and ends with the subject's perceiving.

Therefore

- 7 In all perception, perception takes place after the physical state perceived has ceased to exist.

- 8 The content of perception must be contemporaneous with the perception of which it is the content (i.e. (2) above).

Therefore

- 9 In all perception, the content of perception exists after the physical state perceived has ceased to exist.
10 If one thing exists at a time at which another does not then they are not identical (i.e. (4) above).

Therefore

- 11 In all cases of perception, the content of perception is not identical to the physical state perceived.

The first argument can easily be seen to be valid, so everything turns on the truth of the premises: (4) is an uncontroversial application of Leibnitz's Law, which says that *a* and *b* cannot be identical if one possesses a property—in this case existence at a given time—that the other lacks. Because (1) is, in its natural interpretation, a scientifically demonstrated fact, the pressure is usually put on the philosophical principle (2).

The claim that content and act must be simultaneous is essential to both arguments. I find its intuitive appeal overwhelming, but others do not. The argument for it can be put as follows: (a) The content of a mental state is a—or perhaps *the*—thing that constitutes it. The thought that *p*, for example, is made to be the mental act that it is, in whole or part, by the fact that it is *that p*. The content *qua* content is the core of the act. (b) Anything which wholly or partly constitutes something else must exist at the same time as (or, at least, overlap in time with) the thing it wholly or partly constitutes. Therefore, (c), the content of a perception cannot exist simply at some time earlier than the perception of which it is the content.

The contrary idea, which is that one can perceive directly into the past, involves denying (b). The only half-way plausible way of doing this involves confusing intentional objects *qua* logical feature of a state and intentional object *qua* content. The argument would be that if I think of Charles I, Charles I constitutes the content of that act, but that monarch no longer exists. This is mistaken because the content of the episode of thought is not constituted by the logical object of the thought, but by some vehicle—probably words—by which I think it. I cannot make any sense of the idea that a past object could constitute the

content of a *current* mental state without the aid of a current form of representation.

Given the plausibility of (2), opponents of the argument look again at (1). They argue that the physical object that we see when we look at a star or the sun is not the star or the sun but the light emitted from those celestial bodies. The obvious difficulty for this theory is to answer *where* the light in question is located. If one chose the light that is striking the eye, then one would have the problem of why one did not seem to see two suns, for there are two spatially separate packages of light striking the two eyes: and there might, too, be pressure to treat all visual perception in the same way, and this would be no better than the sense-datum theory. On the other hand, if one chose light at some distance from the eye there would be the problem of deciding how far away the preferred light was deemed to be. How one might specify this distance will emerge in the discussion of the second argument, to which I shall now move.

The second argument's validity is not so intuitively obvious because of the move from (6) to (7). There are two suppressed assumptions, namely:

- 6a In perception the state perceived is the state of the physical world or object as it is at the moment it initiates the causal process that ends in its perception.
- 6b This state can be deemed not to exist at any later time than that at which it initiates the causal process.

(6a) appears fairly uncontroversial, for it would seem that I cannot, at *t*, be perceiving a state of affairs which has not by that time caused any effect in my senses, and, by *t*, no state of the object later than that which initiated the effect on me at *t* has yet affected me (assuming, quite correctly for vision, that it has not been overtaken by a quicker causal process). (6b) seems to be more controversial, because it would seem that, for the distances involved in most visual perception, the object will not have altered its state between the beginning and end of the causal process. One could argue that logically it might have, or one could just individuate states momentarily, so that no state exists *through*, rather than at, a single time. The response to both these considerations is the same. We do not perceive instants, as is shown by the fact that if a change is quick enough we fail to notice it. So, it might be argued, temporal distinctions that are finer than we can notice are irrelevant to the phenomenology of perception.

Travelling at 186,000 miles per second, it takes light $1/186,000$ th of a second to travel one mile and $1/186,000 \times 1/1,760$ th part of a second to travel one yard. Not even the larger of these times is discernible by humans, so it cannot be said that we perceive an object a yard distant as it was $1/327,360,000$ th of a second ago, for such a differentiation does not belong to the world of ordinary experience, but only to the world of science.

On the basis of this response the general argument could be rejected and an answer to the restricted argument propounded, namely that in the case of stars and sun we perceive the light that is that distance away from the eye which is the maximum distance at which the light reaches the eye in a shorter time than the minimum time discrimination we are able to make. This may vary according to the nature and state of the subject. The 'sun' and the 'stars' that we see will, then, be purely visual phenomena, rather like the blue of the sky or a rainbow.

The defender of the time-lag argument might claim that the fact that we cannot discriminate very small time intervals is simply irrelevant. Whether or not I can discriminate such an interval, I do in fact perceive the object a yard away as it was $1/327,360,000$ th of a second ago; and it could, conceivably, have changed in some perceptible respect just in that time, so that by the time I in fact perceive it it is different from how I perceive it to be. The context is relevantly extensional and not intentional: my cognitive abilities in minute time discriminations are not to the point.

The response to this defence would be that it misses the point. In effect, invocation of phenomenal indiscriminability is a challenge to or refinement of the contemporaneity principle that figures in (2) and (8). The point is that the content must be phenomenologically contemporaneous, which means that it must not be possible to discriminate any time distinction between perception and content by direct experience. This condition is satisfied provided that the physical time difference is insufficient to be noticed.

The idea that we can perceive minute distances into the past because we cannot discriminate these times is still open to the attack that it makes no sense for act and content to have a different time, given that the content constitutes the experience. The only way round this objection, if one accepts it, would be to say that these minutely historical events are phenomenologically contemporaneous. This is to admit that there are two different time series, the physical and the

phenomenal, with the latter having a cruder ‘grain’ than the former. It is tempting to think that the very same things could not be in different times in this way, so the phenomenal series would require phenomenal occupants. The only way to avoid postulating separate time series is to take direct perception of the indiscriminably recent past as fundamental.

Assuming that direct perception of the immediately recent past is not impossible, whether this defence against the time-lag argument works depends on two things. On the restricted argument, it depends on whether the naive realist can cope with the purely visual phenomena that it requires or whether they constitute a foothold for the sense-datum theorist. For the generalised argument, it depends on whether the fact that we cannot perceive minute time differences is itself compatible with naive realism. This second point moves us back to the argument from science, for our inability to perceive time in all its detail is only another instance of the way that our perception has a limited resolution and, therefore, does not present things as they are in themselves. Relative to how things are in themselves, our perceptions are, so to speak, somewhat blurred. Whether the naive realist can accommodate facts of this kind is what the argument from science and the argument from illusion dispute.

8 The causal argument

It is an integral part of the atomist picture of the physical world that perception depends on causal chains which link the perceiver and the object they perceive. This argument was, therefore, a commonplace in the seventeenth century. Informally stated, it runs as follows:

There are causal processes running from those objects in the external world that we take ourselves to be perceiving and terminating in our brains/minds (it makes no difference which we say in this context). It is by these processes that our experiences are caused. Experiences include their contents, and so the contents are things caused in us by external objects and are, therefore, different from the external objects we take ourselves to be perceiving.

The reference to experiences including their contents is important, for if one merely concluded that *experiences* were distinct from the objects that cause them, rather than that the *contents* of experiences

were distinct, then the conclusion would be harmless. It is important that the argument establishes that *what I am aware of* is not the external object, rather than that *my being aware of it* is not something external.

Historically, this argument is implicit in the argument from secondary qualities, for if those qualities exist in their phenomenal form only in the senses and as a result of the action of the atomic primary qualities, then the ideas of which we are aware are the product of such a causal chain. The writings of Descartes and Locke are full of remarks that presuppose such a causal perspective. Descartes says:

We clearly see, then, that the sensation of pain is elicited in us merely by the local motion of some parts of our body in contact with some other body; so we may conclude that the nature of our mind is such that it can be subject to all the other sensations merely as a result of other local motions.¹⁹

Locke says, speaking of porphyry,

It has, indeed, such a configuration of particles, both night and day, as are apt, by the rays of light rebounding from some parts of that hard stone, to produce in us the idea of redness, and from others the idea of whiteness.²⁰

The causal account is part of atomism, but that it leads to the refutation of naive realism is more assumed than asserted; for, if it is taken for granted that Descartes' sensations and Locke's ideas lack intentionality, then the fact that these are the things in the mind in perception is enough to refute naive realism.

Causal argument formally stated

Putting the causal argument more formally, one might start with a version that is too weak:

- 1 Perceptual experience occurs at the end of a causal chain running from the external object to the perceiver's brain/mind.
- 2 Things located at the opposite ends of a causal chain cannot be identical.

Therefore

Perception

- 3 Perceptual experience is not identical with the external object perceived.

The argument is sound but harmless, for no one wishes to identify the object with the experience, *simpliciter*. So one might instead say that it is the content that is caused:

- 4 The content of experience occurs at the end of a causal chain running from the external object to the perceiver's brain.
- 5 Things located at the opposite ends of a causal chain cannot be identical.

Therefore

- 6 The content of experience is not identical with the external object perceived.

The argument is again valid, but it can be challenged either by denying (4) or by denying the significance of (6). A naive realist would argue that the content of experience is the object perceived—e.g. a table or a man—and this is not caused by the process involved in perception. (1) is true but (4) is false and rests on a confusion of the act and the object of experience. As the argument should start from the undisputed fact, the serious version runs as follows:

- 7 Perceptual experience occurs at the end of a causal chain running from the external object to the brain.
- 8 The content of perceptual experience is part of, or internal to, the experience.
- 9 If something occurs at a certain location anything which is part of or internal to it occurs at (or within) that location.

Therefore

- 10 The content of perceptual experience occurs at the end of a causal chain running from the external object to the brain.

The argument is then completed exactly as before:

- 11 Things located at opposite ends of a causal chain cannot be identified.

Therefore

- 12 The content of perceptual experience is not identical with the external object perceived.

The task of defending this argument consists in proving that (8) is correct under an interpretation of 'content' strong enough to make the conclusion significant. Part of this task can be seen as defending (9), for one excessively weak sense of 'content' is one which makes it synonymous with 'intentional object'. It could be argued that Charles I is internal to or part of my thought that Charles I was a martyr, but that he does not occur at the same location as my thought but 340 years prior to it (or, alternatively, in heaven whilst my thought is mundane).

Proving that experience and its content cannot be separated is, in effect, the point of the argument in Chapter VI. That argument is constructed from an amalgamation of the causal argument and the argument from hallucination. And proving that the nature of the content which is inseparable from experience is substantial enough to be inconsistent with naive realism—that is, to amount to a sense-datum—is the purpose of Chapter VII. Further development of this argument, therefore, must await those later chapters.

9 The argument from hallucination

The fact that people have experiences which are somewhat like perceptual experiences but the content of which is certainly 'in the mind' has always been used to support sceptical conclusions. Descartes' appeal to dreaming is an instance of this.²¹ The strategy in this form belongs to the classical sceptical tradition that I outlined in Chapter II. It is used to show, not that what we are aware of is subjective and mental, but that we can have no reasonable confidence that, in a given case, it is the real world that we are seeing. As an argument against naive realism, it is, I think, modern. Taken on its own it is not a strong argument. It is seen most easily set out formally:

- 1 The contents of hallucinations are qualitatively similar to those of perceptual experiences.
- 2 The contents of hallucinations are subjective images or 'sense-data'.

Therefore

- 3 Perceptual experiences have as their content subjective images or 'sense-data'.

Various quibbles can be raised against this argument. It has been pointed out that hallucinations are not actually very like perceptions; they are chaotic and fantastic. This is doubly irrelevant. It is irrelevant because there surely could be hallucinations which were indistinguishable from veridical perceptions. And it is irrelevant because all the argument as it stands requires is a general qualitative similarity between perceptual and hallucinatory contents. It does not matter, that is, that the pink rat suddenly turns into Mrs Thatcher; all that matters is that there seems to be something *pink* and of a shape of a sort something could really look to be. Nor is there force behind the objection that hallucinations are not experiences with content because they only *seem to be* experiences and seem to have content: or that one only *believes* one is having an experience when one is not. For it is possible to be quite aware that one is hallucinating and hence to have no false beliefs about the situation; and there is an introspectable content to the experience which cannot be identified with any external state of affairs, so it is quite innocuous to characterise it as a subjective image.

The substantial objection to the argument is that it is invalid. In order to make it valid requires some such premise as:

- 2a Things can be qualitatively similar only if they are of the same general kind and have the same ontological status.

This has a certain intuitive appeal—how can things really be of a radically different kind if they share properties? The example of forgeries is often used against it: a forged and a genuine banknote are qualitatively exactly similar, but one is a *real* banknote and the other is not. This argument is not very powerful because a forged note and a real one are both straightforwardly physical objects—the differences are a matter of origin and convention, and not of natural ontological category. Nevertheless the principle is not so obviously true that it can be used to support a controversial conclusion. Berkeley's principle that 'an idea can be like nothing but an idea' is a more specific version of (2a) and could be expressed as:

- 2b Something can be qualitatively similar to a mental item only if it is a mental item itself.

Once again, this is question-begging, but one has the feeling that there may be something in it. Everyone agrees that some things, such as pains and tickles, could not be copied outside the mind.

Berkeley argued that all sensory 'ideas' were inseparably linked to such sensations. His arguments have not convinced, but there remains the puzzle of why pains and colours should be so different. As far as I understand the 'adverbial' theory of sensory content, they are not different according to that theory because pain and phenomenal colour are both properties of mental activities, and it is very plausible to claim that nothing could be like a feature of a mental act except a feature of a mental act.

The most common way of distinguishing sensations from perceptual content is to claim that the latter consists of intentional objects and that intentional objects do not actually instantiate the sensible qualities that they apparently present. Thus when the drunkard hallucinates a pink rat there is nothing pink; the drunkard's mental image, whatever that may be, is not a pink thing. On the other hand, pains are genuinely painful and itches, itchy. Once again we are back to the intentionality of perceptual content.

But even if one were to accept the view that the hallucinator had a pink sense-datum, the claim that both sense-data and physical objects could be coloured and rat-shaped would not be manifestly false. It may be odd to think that both a mental object and a physical one could share a type of intrinsic property, but no real reason has been given for thinking it impossible. The argument from hallucination is, therefore, question-begging.

10 Conclusion of Chapters I-III

None of the arguments in these chapters is conclusive; but the intentionalist response to illusion looks, in my view, fragile and is still to be properly investigated: the argument from secondary qualities, I concluded above, is persuasive, if not definitive. Responses to the argument from science are inadequate, except for those that rely heavily on intentionality. The time-lag argument can be avoided for most phenomena, but it is not clear that the response is convincing. To this point, we seem to be in a 'seven leaky buckets' predicament: several arguments that are not individually conclusive. The problem is to decide whether they add up to a good cumulative case and hold water, or whether as a group they are no better than the best alone. Are they unsuccessful deductive arguments or mutually supporting evidence for a certain way of understanding perception? Fortunately we do not need to decide this. As so far stated, the causal argument and the

Perception

argument from hallucination have little force, but they can be brought together to make a powerful argument for sense-data. This, however, must wait until Chapter VI. The next two chapters will be concerned with two major planks in the modern alternative to the sense-datum theory. Chapter IV will deal with—and, I hope, dispose of—the anti-private language argument, and Chapter V will attempt to show the inadequacy of the physicalist theories of perception that are the fashionable alternatives to the sense-datum theory.

CHAPTER IV

Sense-data and the Anti-private Language Argument

1 Varieties of the argument

We have seen how the traditional arguments for sense-data can be resisted. The supposed *coup de grâce* against the theory is provided by an attempt to prove that, whatever the arguments for sense-data, there just could not be such things. Sense-data, at least as usually conceived, are not public objects in public space, but are necessarily private to individual observers. They are what have come to be called *logically private objects*. The final crisis for the empiricist conception of perception was precipitated by Wittgenstein's famous polemic against such private objects. He argued that such things could not be objects of linguistic reference and, hence, not objects of thought or consciousness.

Many uncertainties attach to Wittgenstein's argument. Most importantly, it is disputed how the argument actually goes and what it is meant to prove. There is a common idea behind all versions of the argument. This is that there could be no language to describe logically private objects because the kinds of constraints that make for stability of use—and therefore for meaningfulness—are essentially public. Two different kinds of argument, representing different interpretations of Wittgenstein, have been brought to articulate this intuition. One—the more traditional interpretation—emphasises the impossibility of verifying whether use is consistent when the standard of correctness is given by something which is logically private. The second is a form of radical nominalism. Consistency of use for a referential term involves referring to *the same kind* of thing. Wittgenstein is interpreted as arguing that what counts as 'the same', whether public or private, is no more than what the linguistic community agrees to call the same, and the community

would have no access to the logically private. There is also a third strategy available. This is to abandon the attempt to argue directly for the conclusion of the anti-private language argument, and instead to present it as part of a radically alternative picture of how language works which we should adopt because it avoids certain traditional problems in epistemology and the philosophy of mind.

Discussion of privacy is complicated by the intermeshing of Wittgensteinian exegesis with the question of philosophical truth *per se*. Of course, it is the philosophical question of whether there can be logically private objects that is the major question, but, in practice, Wittgenstein's role in the argument is so central that questions of interpretation naturally occur in the course of the discussion.

Before considering the structure of the argument I shall look at what it is meant to prove.

2 What the argument is meant to prove

There are two views about what Wittgenstein's argument is meant to prove. On the first and stronger interpretation the argument is intended to show that there are no logically private objects and hence no special, private and subjective dimension to human experience. If this interpretation is correct Wittgenstein is a behaviourist or perhaps some other sort of hardline physicalist because he is denying the existence of a private subjective realm. According to the second and weaker interpretation, Wittgenstein is not concerned to deny the reality of logically private forms of awareness, but only to affirm that the contents of the subjective dimension have their form dictated by public criteria. This is not a denial of the logically private, but an affirmation of the dependence of the private on the public. On this interpretation the argument is less disturbingly reductive.

It is possible to argue either about which of these two views is Wittgenstein's, or, more generally, about whether an argument of the same general sort as Wittgenstein's can be provided for the weaker conclusion. The availability of the weaker option seems more controversial than the availability of the stronger: the *prima facie* availability of behaviourism—as opposed to whether it is true of conscious beings—is not at stake; it is not so clear that one can have some of the advantages of behaviourism without the implausible consequence of denying altogether the reality of subjectivity.

I shall briefly present three arguments for thinking that Wittgenstein took the hardline view, and then investigate whether, on quite general grounds, the weaker compromise position is tenable.

1 Wittgenstein's treatment of first-person reports of experience tells strongly for the reductive interpretation. Wittgenstein denies that an utterance such as 'I am in pain' has a truth-value.¹ But if there were a private sensation at all, albeit one which merely reflected the public criteria by which someone could assess such an utterance, it is difficult to see how the utterance could fail to be either true to or not true to that sensation. If there is something at all for the utterance to be about which constitutes a private sensation then the statement should have a truth-value.

2 In *Philosophical Investigations* II, iv, Wittgenstein writes:

My attitude towards him is an attitude towards a soul. I am not of the *opinion* that he has a soul.

When put together with (1) this strongly suggests that 'having a soul' (that is, being the sort of creature to which typically human and psychological predicates attach) is not a matter of fact but of value; it is a question of how an organism is regarded. The theory of mind implicit here seems to be parallel to an emotivist theory of ethics. This at least would save Wittgenstein from the paradoxical position implied by (1) that there is an asymmetry between first-person and third-person uses of sensation ascriptions such that there is nothing to make the former true, but there is something which makes the latter true. If we take the 'emotivist' view of psychology there is nothing which makes either true. There is a difference, as Wittgenstein plainly insists, but that consists in the fact that there are criteria for assertion (which are not the same as truth conditions) in the third-person case, but not for the first person.

3 In the famous 'beetle in the box' analogy Wittgenstein claims that 'the box might even be empty'.² If this is so, then no logically private element can be involved in the ascription of sensations. He says:

The thing in the box has no place in the language-game at all; not even as a *something*: for the box might even be empty.

If the weak interpretation were correct there would be a *something*, though *what it was* would be determined by external criteria.

Whether or not the weak theory is Wittgenstein's, it has more appeal than the strong in its own right. Can it be defended? One could argue

that it cannot be defended by trying to put the protagonist of the weak theory onto the horns of a dilemma.

We ask him whether the subject can recognise the private component in his experience, so that he can tell when it recurs. If he were to answer 'yes' to this, then it is unclear why the subject cannot engage in the exercise Wittgenstein debunks in *Philosophical Investigations*, I, 258, of writing down a name for the logically private component whenever it occurs and hence developing a private language. If, on the other hand, he denies that the subject can recognise the private component, then it would appear that the whole point of allowing such a component has disappeared. The reason for preferring the weak to the strong theory was that it permitted a real, subjective element, which gave body to experience; but if that real subjective element is such that we cannot tell directly the difference between, for example, the sort of subjective element that accompanies an intensely painful burn and the sort that accompanies tasting claret, then it hardly constitutes what we think of as the subjective dimension of private experience. Wittgenstein would be quite correct in maintaining that in this case the box might as well be empty.

There is, however, a way out of this dilemma. The weak theorist can say that we do recognise the private experience, but not in virtue of any in-principle private feature of it; rather we recognise it in terms of its apparent public object. So I recognise the experience as being one as of seeming to see something red, for example, where 'red' is understood as a public property of public objects. The experience is logically private to me, but its kind of content is not. So if I write down the name for the object of experience it will not be some new name in a private language but a public one that is already in currency.

Whether this strategy is available will depend on the kind of experience one considers. Sensations, as opposed to apparent perceptions, do not fit in so easily. The qualitative nature of a pain or an itch is not naturally attributable to a public object, in the way that redness is. One way round this is to characterise sensations, not via their objects, but via their causes and/or their effects. In this case one thinks of a pain as *what goes on in someone when they have been damaged or when they are disposed to aversion behaviour*. This is a very counter-intuitive theory of sensation: it is the topic-neutral analysis provided by mind-brain identity theorists, such as Smart and Armstrong, and it says that we have no direct knowledge of the

intrinsic qualitative nature of sensations.³ At least, in the case of pain, the association with particular stimuli and responses is clear and is part of the way we normally think of pain, but what the input/output conditions are for a slight dull ache is anyone's guess; they certainly don't seem to define what the ache feels like.

Even in the perceptual case, there are problems with the primacy of the public. It presupposes, of course, that an objectivist account of secondary qualities is plausible, but it also leaves a problem for phenomena that are not like any physical state. One might, for example, hold that though something can look blurred, nothing can actually *be* blurred. For our purposes, however, the main point is that this weak theory leads back to problems with naive or direct realist theories of perception. It suggests that when someone seems to see red, but there is nothing red present, either because the object seen is some other colour or because the experience is hallucinatory, then what there really is is an experience *as of* an apparently public object. This still leaves us with a need for an analysis of such experiential states. The stronger position just denies that there is anything to report in such cases, but, once this view has been abandoned, then an account of what is reported is required. I discuss the 'as of' theory of the objects of experience in Chapter VII and find it wanting. Perhaps some other way of characterising the content of illusions and hallucinations could be found for the weak position; but they would all have to be in terms of the apparent presence of a public object and so they would have to employ the notion of an intentional object. Chapter VII tries to take on all the likely variants of the intentional object theory. If, as I believe, they are all inadequate, so is the weak interpretation of Wittgenstein.

3 The traditional interpretation of the anti-private language argument

We can put aside for the moment the question of what the argument is best understood as trying to prove, and consider what it is in itself. According to the more traditional interpretation, the conclusion follows from Wittgenstein's philosophy of language *augmented by certain considerations that apply only in the case of discourse concerning the logically private*. According to Kripke's more recent interpretation, on the other hand, it follows directly from Wittgenstein's philosophy of language alone.

As I shall reconstruct the traditional interpretation it begins with two premises derived from Wittgensteinian philosophy of language in general and is followed by a third relating only to private objects. It is this last which bears the main burden of the argument.

- 1 The meaning of a word is given by the rule which governs its use.
- 2 A rule must be such that there is a difference between following it correctly and not following it correctly.
- 3 If a word purports to name a logically private object then there is no difference between following the corresponding rule correctly and not following it correctly.

Therefore (from (2) and (3))

- 4 If a word purports to name a logically private object it does not follow a genuine rule.

Therefore (from (1) and (4))

- 5 If a word purports to name a logically private object then it lacks meaning.

Although philosophers have managed to argue about whether language is 'rule-governed', premise (1) is usually taken to be uncontroversial. There appears to be a sense in which it is platitudinous, because you cannot use words just as you like if you are purporting to speak a meaningful language. Premise (2) seems only to fill this out: if a rule does not exclude some things and include others it has no content. The argument is valid, so the weight of a challenge to its soundness will fall on the truth of premise (3). As it stands (3) is a bare assertion and so everything turns on producing an argument for (3). The argument for (3) is drawn from the famous section 258 of *Philosophical Investigations*:

Let us imagine the following case. I want to keep a diary about the recurrence of a certain sensation. To this end I associate it with the sign 'S' and write this sign in a calendar for every day on which I have the sensation.—I will remark first of all that a definition of the sign cannot be formulated.—But still I can give myself a kind of ostensive definition.—How? Can I point to the sensation? Not in the ordinary sense. But I speak, or write the sign down, and at the same time I concentrate my attention on the

sensation—and so, as it were, point to it inwardly.—But what is this ceremony for? for that is all it seems to be!

A definition surely serves to establish the meaning of a sign.—Well that is done precisely by the concentrating of my attention; for in this way I impress on myself the connection between the sign and the sensation.—But ‘I impress on myself can only mean: this process brings it about that I remember the connection *right* in the future. But in the present case I have no criterion of correctness. One would like to say: whatever is going to seem right to me is right. And that only means that here we can’t talk about ‘right’.

The last two sentences of this section give us two premises which together entail (3).

- 6 If a word purports to name a logically private object then there is no difference between correctly using it and only seeming to use it correctly.
- 7 If there is no difference between using a word correctly and only seeming to use it correctly then there is no difference between following the corresponding rule correctly and not following it correctly.

The controversial premise here is (6). One way of understanding Wittgenstein’s argument is to see it as a *reductio ad absurdum* of the hypothesis that there are, or can be, words which derive their meaning by designating logically private objects. The reduction to absurdity consists in showing that the hypothesised private discourse could not meet certain constraints which any possible form of language must meet and would, therefore, not be a genuine form of discourse at all. For the argument to possess any persuasive force it must not rest on any premises which the private linguist (or anyone tempted to be a private linguist) would, by dint of being a private linguist, refuse to accept, unless that premise can be independently demonstrated to be true. As I have so far developed the argument, (6) is a premise for which no independent reasons have been given, and a little reflection is sufficient to show that (6) is a premise that the private linguist has every reason to reject. The denial of (6) follows from the hypothesis that there are logically private objects for the following straightforward reason.

One thing normally involved in calling something an object, is that the thing possesses some particular determinate nature. In calling something a logically private object it is intended that it possess some definite qualitative nature—it is something that is as a matter of fact thus and so. I take it that this is something to which anyone believing in logically private objects will assent. It follows from the fact that this is what a logically private object is, that there must be a difference between applying the name of a logically private object correctly and only seeming to do so. If we take as an assumption of the anti-Wittgensteinian that there is some private object *O* with a determinate nature *N* occurrent at time *t*, then that there was *O* with *N* at *t* is a statement true at all times after *t*. If the subject *S* who has *O* coins the name ‘*A*’ as a general name for all *O*s of the type *N* and proceeds to apply ‘*A*’ to some other *O* at a time later than *t*, then this either will be, or will not be, of type *N*; if it is of type *N*, then *S* uses ‘*A*’ correctly, and if not, *S* does not. In other words, given the supposition that there are determinate private objects of experience, it follows that such objects of given types can be re-identified correctly or incorrectly in just the same sense as applies to other determinate objects.

James Hopkins calls the conception of private objects contained in this argument the ‘recognitional conception’, for it rests on the idea that if there were private objects they would possess determinate natures waiting to be recognised.⁴ Wittgensteinians must attack the recognitional conception. They can either attack its application to private objects—that is, they could deny that private objects could have a nature waiting to be recognised, whilst allowing that public objects can have such a nature: or they could attack it generally—that is, they could attack the notion that features of the world, private or public, are objectively present waiting to be recognised, existing prior to and independently of the development of our concepts. This latter position would involve a strong form of nominalism, for it involves denying that there is any sort of objective grounding for our predicates. The traditional interpretation of the argument involves the more restricted programme of attacking the application of the recognitional conception to private objects. Kripke’s interpretation involves the full anti-realism of the second strategy.

The most usual versions of the traditional argument import what look like forms of the verification principle into the argument.

4 The verificationist attack on the recognitional conception

Usually when interpreting the argument philosophers light upon the sentence in 258 which precedes the last two, which I turned into (6) and (7). They are impressed by the statement that ‘in the present case I have no criterion of correctness’. The concept of a *criterion* in Wittgenstein’s philosophy has elicited an enormous amount of discussion. It seems to be generally accepted that a criterion is not a truth condition (if it were it would not be inconsistent with the recognitional conception, for the nature of the object itself would then be the criterion for the application of its name) but is a necessarily reliable (though not infallible) way of telling whether a concept applies. Because it relates to *telling* whether a concept applies, it has verificationist overtones. The recognitional conception of private objects challenged line (6) of the original argument. We need a version of (6) which paves the way to rule out the recognitional conception. This will be done if we can tie the notion of how something objectively is to whether we can tell how it is. Thus we have

- 6 If a word purports to name a logically private object then it is not possible to tell whether one is using it correctly or only seeming to use it correctly.
- 8 If it is not possible to tell whether one is using a rule correctly or only seeming to do so, then there is no difference between using the word correctly and not doing so.

It is (8) which makes the apparently verificationist connection between objective correctness and being able to tell whether one is using a word correctly. From (6') and (8), (3) follows directly by hypothetical syllogism.

The trouble with the argument now developed is not directly that it imports verificationism. Language is something we use and the suggestion that meaningfulness requires that we be able to tell whether we are using the words properly, rather than that there be some entirely objective correspondence between the word and the world, is not at all implausible. The trouble with (6') and (8) is that it is not clear what ought to count as being *able to tell* whether one is using a word properly. The same ambiguities infect ‘being able to tell’ as infect verificationism. Is ‘being able to tell whether p’ like weak verification of p—that is, does it simply require that we have

some evidence: or is it the strong verification—that is, does it require evidence amounting to conclusive proof? The latter is so strong a requirement that nothing empirical will meet it, the former is so weak that ‘seeming to remember’ how one previously used a word should pass the test. Much of the literature on the private language argument has concerned whether logically private memory—that is, memory for the accuracy of which there can be no direct, public test—counts as a means of telling whether a word is being used correctly. It certainly does not constitute strong verification, but it seems to constitute weak verification, for how things seem to me is, in most contexts, *some* evidence of how they are.

It is possible to construct middle positions between weak and strong verificationism, but none of them seem to have anything to recommend them except that, if correct, they entail what we are trying to prove, namely the primacy of the public; they are not fitted to function as arguments for that position. For example, it might be thought that what matters is that a statement should be verifiable publicly; but this simply appeals to the same intuition as does the anti-private language argument. As it is the priority of the public that we are trying to justify, it will not do to use such a principle to justify the argument.

A somewhat different principle is that any statement must be testable in at least two ways: however, again this seems to lack sufficient reason, possessing as its rationale simply that it grants what is wanted, but it does not justify it. A less arbitrary version of the same principle would be to insist that an undeterminable list of ways of testing any statement must exist, such that the process of verification never has to come to an end. However, if this means that there must be an inexhaustible supply of different *ways* of verifying a statement—i.e. each one employing an in-principle different method, like using a different sense—then no statement is verifiable in this sense; there are limits to the types of verification available. If, on the other hand, it means that there must always be the possibility of a further test, albeit of a type already employed, then it seems to me that such an inexhaustible supply is available to test claims about private objects, for it is always possible to try again, or try harder, to remember. Wittgenstein scorns this method, comparing it to buying a second copy of the same newspaper to check on the news in the first copy, but this is an unsound objection at this juncture.⁵ We do in normal cases exhort people to try harder to remember to see whether they can improve on

their first memory attempts; if this is not to count in the present case, that can only be because we already have grounds for discrimination against memory of logically private objects; which is the matter at issue.

It seems, then, that neither strong nor weak verification principles will do the job, and though middle positions are conceivable, they do not function as support for discriminating against logically private memory, for that they do so discriminate seems to be the only reason for adopting them, as they have no independent plausibility. It seems to follow that if there are to be sound objections to allowing private memory to count as a means of telling whether or not one is following a rule correctly, they will not be based on some independently plausible verification principle. The argument against private experience as a means of verification will have to be independent of more general considerations about the verification principle.

Wittgenstein's argument for (3) in fact rests on the following two propositions:

- 6" If a word purports to name a logically private object then the only criterion on which its use could rest would be how it seems to the subject.
- 9 If the only criterion for the use of a word was how it seemed to the subject, then there would be no difference between its being used correctly and its not being used correctly.

(6") and (9) together entail (3).

The major question for the traditional interpretation is how one might defend (9). One path is to argue that (9) is not just a principle about the use of language, but is a special case of a general principle:

- 10 If the only ground for a belief or judgement was, necessarily, how it seemed to the subject then there would be no difference between that belief's being true and its being false; that is, the question of its truth-value would not arise.

If one accepts (10) then a principle concerning the conditions for word-meaning is being backed up by one that concerns judgements in general. It follows from (10) that a particular uncheckable judgement—such as that I seem now to see red—would have no truth-value, even though there was nothing contentious about any of the terms used. The question of whether (9) rests on (10) is intimately connected with whether one makes the strong or the weak interpretation of what the

argument is meant to prove. For even if a subject could report their experience only using terms with public reference (as in 'I seem to see red'), nevertheless it would still seem to be a logically private matter whether that was indeed how it seemed to them, if there is a fact of the matter, as the weak interpretation requires that there be.

Two arguments can be brought against using (10) as the grounds for (9). First, it is question-begging. The argument as so far developed runs from (1) to (5) with (3) the crux. (3) is supported by (6") and (9):(6") is not controversial and (9) follows from (10). (10) is, therefore, the central point of the argument. But (10) will no more be accepted by someone who believes in the recognitional conception of private objects (and all believers in private objects should accept this conception) than would the original (6): if there really is a fact to be recognised about how it seems, then there must be a fact about whether it has been correctly recognised. It appears that on this interpretation of the argument it is question-begging, resting on an assumption that anyone who believed in private objects would reject.

The second objection against (10) has, in effect, been presented by A.J.Ayer.⁶ Ayer points out that, as our knowledge of the public world depends on experience, it must, therefore, depend on how things seem to us as individual subjects. If our conception of the public world depends on our ability to recognise when things look to us similar and when different, and in what respects, then it cannot be the case that there is truth and correctness in our judgements about the former if there is not truth and correctness in our judgements about the latter. The burden of Ayer's point could be put by saying that how something *public* seems to me is just as much dependent on my uncheckable, criteria-less subjective judgement as is my judgement about my private sensations, and my grasp on the world depends on how I experience it—that is, on how it seems to me.

If Ayer is correct, the argument will move from being an attack on private *objects* of experience to being an attack on the privacy of experience in general. Ayer's opponents presumably do not want to deny that our conception of how the world is in some sense depends on experience, so they will affirm the two following propositions:

- a Our conception of how the world is rests upon our experience of it.
- b Our conception of the world does not rest on how it seems to us, in that sense of 'seems' according to which there are no

further criteria for how things seem other than the subject's own judgement.

There follows from this

- c Experience is not identical with how things seem to the subject, in that sense of 'seems' in which there are no further criteria.

There are various ways of trying to make this conclusion acceptable. One is to argue that experience is not a sort of intermediary between the subject and the world, and that to interpret 'seems' statements as reports of experience is to treat it as such. Hacker says that 'seems' statements are not reports, but 'display extensive affinities to the kinds of utterance, which, in the wake of Wittgenstein, have been called "avowals"',⁷ and that, *a fortiori*, they do not report anything that constitutes evidence for claims about the world. Reports of experience express 'how something perceptually strikes one'. I have argued above (in Chapter II) that this approach is necessarily committed to treating 'seems' judgements as merely tentative expressions of belief, and that this leads to a 'belief analysis of perception. I also argued that treating *seemings* as tentative beliefs is wrong, and shall be arguing against the 'belief' analysis of perception in the next chapter. This route out of Ayer's objection is, therefore, not available.

There is another line of response to Ayer's argument that is tempting but rests on a mistake. It might seem that because Ayer says we can recognise the world only if we can recognise what it seems like, then we must have concepts that relate to subjective experience before we can have ones that relate to the physical world. This is taken to be equivalent to saying that we must be conceptually aware of experience *as such* in order to be aware of the world. It is fairly obvious that we do not move from conceptualising experience to the world in this way, and even more obvious that other animals do not do so: yet we and the animals depend on experience to find out about the world, so such dependence does not require a conceptual grasp on how it seems, as the argument claims. This response misses the point. It will not suffice for the Wittgensteinian to claim that our grasp of concepts which characterise the public world is prior to our grasp of concepts characterising experience. Ayer's point—that our grasp on how the world is depends on how it seems to us—neither says nor entails anything about our having a concept of experience. The world can

seem thus and so to us without our grasping it as seeming (as opposed to being) thus and so. Ayer is only claiming that our conception of the world depends on our experience of it, not that it depends on our conception of our experience of it. It is enough that we are *in fact* judging on the basis of how the world seems to us, and that such 'seeming' is, in fact, logically private.

Resistance to Ayer's argument stems from a feeling that Ayer is reifying experience and treating that into which it is reified as an object of experience. This, in general, is what people fear the sense-datum approach does. In fact Ayer's argument involves no such reification and is neutral about how one should analyse experience. It requires only that recognition of features of the world depends on the subject's experience and that, in the last instance, a subject cannot have and does not need public criteria for how it seems to them in such experience. All such experience might be in the form of being *as of* a physical world, and the subject may or may not have the concept of experience.

Ayer's argument seems to me to prove the falsehood of (10) and, therefore, that (9) cannot be justified on the basis of (10). In fact, I think (9) is meant to be very close to being a trivial analytic truth, dependent on no further principles. It is meant, I think, to follow from the idea that where there is objective truth one cannot say just what one likes, because there has to be a logical distinction between a judgement and its object—that is, between a judgement and what makes it true or false. The transition of ideas is as follows. If a judgement were correct just in virtue of its being judged, then there would not be the right sort of distinction between judgement and object: reality would, so to speak, be following belief, rather than vice versa. Then there is the thought that something's seeming to be the case is an instance of judging it to be the case, so if seeming guarantees truth then there is no gap between judgement and object ('what makes it true') and so no truth.

The platitude behind all this is that one does not make something true by judging it to be so, in a sense which would make reality follow whatever happened to be the judgement. But this is not the situation when 'how it seems' is the only criterion of truth. First, the near-infallibility of certain judgements about one's own experience is not a matter of the judgement establishing the reality, but of the lack of room for error in judgement, given the kind of reality we are talking about, namely, one with which one is directly acquainted. So

the logical distinction of judgement and object and the priority of the latter over the former is not upset. Second, there is not an absolute infallibility, as Wittgenstein's argument presupposes, but only the kind of privileged access that makes certain kinds of mistake virtually unintelligible in a sane person.⁸ Third, there is an equivocation in the sense of 'seems'. One can interpret the expression 'whatever is going to seem right to me is right' as meaning 'necessarily, whatever I judge to be right is right'. This is the sense Wittgenstein requires. But if one interprets 'seems' in its phenomenal sense it just means 'whatever phenomenally seems to be the case is how it in fact is, experientially'. This is a harmless tautology, for the expressions 'how it phenomenally seems' and 'how it experientially is' are, in one of their senses at least, synonymous. The question of how the experience is, is not a matter of how it judgementally seems to me, though it does have consequences for judgement, for there ought to be no barrier to my correctly recognising how it seems, in certain fundamental respects, at least.

Finally, there must be something wrong with these Wittgensteinian reflections on 'seeming' for they entail (10), which has already been refuted. They entail (10) because, if they were correct they would work against all necessarily uncheckable seemings. In all such cases, its seeming to be so is the only criterion for correctness, and if the distinction between seeming and being cannot be made under these circumstances then all judgements necessarily dependent on how it seems to the subject will be vacuous.

It seems that there is no way of grounding the traditional form of the anti-private language argument.

5 Kripke's radical interpretation: the outline

The crucial difference between Kripke's interpretation of the argument and the traditional one is that the object of the attack is not *rules which purport to relate to logically private objects*, but *the private understanding of rules as such*, whatever their subject-matter.⁹ The scope of the argument is, therefore, enlarged in just the sort of way that Ayer's criticism enlarged it, when he showed that all judgements of how things seem to the subject came under Wittgenstein's hammer, irrespective of whether the ultimate object of the seeming was private or not.

The argument can be reconstructed in the following way. The first three steps concern general features of meaning and rule-following. The first premise is the same as the first premise of the traditional argument, the second is an explanation of what is involved in following a rule and the third follows from the first two.

- 1 The meaning of a word is determined by a rule.
- 2 A rule is something which determines that a word is used in a consistent manner: following a rule involves using a word 'in the same way' and recognising as 'the same' the circumstances which are the criteria for its use, which will often mean recognising something as being of *the same kind of thing* as the kind of thing to which it is supposed to refer.

Therefore

- 3 Meaning requires that there be a standard for 'doing the same thing' and for being *the same kind of thing*.

The second stage in the argument is the consideration of the suggestion that our understanding of rules is something internal, mental or private. Wittgenstein considers three ways in which this might be so:

- 4 If something internal or mental constituted our understanding of a rule for the use of a word then that understanding is either (a) a picture of what we are to recognise as falling under the word; or (b) a *disposition* to 'do the same thing' in the way appropriate in our use of a particular word; or (c) some intuitive grasp on what constitutes 'doing the same thing' in the way appropriate for a given word.

Using arguments we shall consider below, he dismisses all three.

- 5 Our understanding is not a picture which we apply.
- 6 Our understanding is not a disposition.
- 7 The contents of our minds are finite, so nothing that we intuitively grasp can determine the whole content of a recursive rule and hence cannot determine what it is to 'do the same thing' open-endedly as language competence requires.

Therefore

- 8 Understanding a rule does not consist in something internal or mental.

I shall consider below the arguments for (5), (6) and (7). Wittgenstein's reason for dismissing (5) is fairly uncontroversial. His arguments for (6) are very weak, but they have been augmented by Kripke. (7) is the crux of the Kripkean reconstruction.

Having dismissed the mentalistic account of understanding Wittgenstein considers the alternative.

- 9 As meaning requires that there be a standard for 'doing the same thing' (i.e. given (3)) and as this is not internal (i.e. given (8)) then the standard for 'doing the same thing' must be constituted by social practice.

Therefore

- 10 The correctness or incorrectness of all judgements of similarity is a function of social agreement.

Finally, these conclusions about recognising what constitutes being the same, or doing the same thing are applied to logically private languages in a very direct way.

- 11 If we could name or recognise logically private objects, judgements of similarity concerning them would not be based on social agreement, but on uncheckable private judgement.

Therefore

- 12 There are no logically private objects which we can name or recognise.

6 The problem with the infinite content of a rule: Wittgenstein's treatment

The problem of how a mind can grasp the recursive content of a rule—that is, the theory of understanding contained in (4) above—is plausibly described in *Philosophical Investigations* 139:

When someone says the word 'cube' to me, for example, I know what it means. But can the whole *use* of the word come before my mind, when I *understand* it in this way?

Well, but on the other hand isn't the meaning of the word also determined by this use? And can't these ways of determining meaning conflict? Can what we grasp *in a flash* accord with a

use, fit or fail to fit it? And how can what is present to us in an instant, what comes before our mind in an instant, fit a *use*?

Wittgenstein considers various responses to the problem which are realist in the sense that they all entail that there is something we grasp which determines the use. The first suggestion which appears above as (4a), is that understanding the word 'cube' involves having a picture of a cube in the mind's eye, which illustrates how the word is to be applied. He correctly rejects this on the grounds that one needs to interpret the picture to know how to apply it.¹⁰ Hence the conclusion (5). The more important suggestion is that a concept (or the possession of a concept) is a disposition. This is important because realists about concept possession, whether they are mentalists or physicalists, tend to treat concepts as fundamentally expressed in dispositions or capacities. This is most clearly the case for those physicalists who treat the possession of a concept as a disposition to linguistic and other behaviour housed or 'realised' in some neural machinery. Dualists may be tempted to treat the possession of a concept as essentially contemplative, involving the intellectual apprehension of a universal, but nevertheless, this is almost completely expressed in terms of certain capacities. Wittgenstein's treatment of the dispositional theory is brief and inadequate. It occurs in paragraph 149:

If one says that knowing the ABC is a state of mind, one is thinking of the state of a mental apparatus (perhaps of the brain) by means of which we explain the *manifestations* of that knowledge. Such a state is called a disposition. But there are objections to speaking of a state of mind here, in as much as there ought to be two different criteria for such a state: a knowledge of the construction of the apparatus, quite apart from what it does.

Wittgenstein here makes two claims. First that a disposition involves a categorical base and second that knowledge of a disposition (criteria for the ascription of a disposition) involves knowledge of that base as well as of the type of activity involved in the disposition. The first claim rules out by fiat two conceptions of disposition. First it rules out what has been called the 'phenomenalist' account of dispositions, as employed, for example, by Ryle.¹¹ According to that theory, to ascribe a disposition is to attribute certain performances from the past

and to predict that they would be repeated under appropriate circumstances in the future. 'Disposition' in this case does not refer to or involve something actual which causes or otherwise underlies the performances. This theory would be entirely unsuitable for explaining what someone grasps when they grasp a meaning, because it allows for nothing but the performances. The terms of reference of 149 would require a realist, not phenomenalist account of dispositions. But there is a realist theory which does not involve a categorical base, namely one which treats a mental capacity as an irreducible real power. Many philosophers and scientists treat the basic forces in nature realistically, but not as expressions of some lower level of structure. There seems to be no *a priori* reason why mental powers could not be treated similarly.

Wittgenstein's second point omits what would now be characterised as a topic-neutral account of our knowledge of the categorical base for dispositions. A physicalist exponent of the causal theory of mind would hold that a concept was a disposition with a neural basis, but that all that we know directly is the content of the disposition: we know that it must have a basis, but we know that only by the topic-neutral characterisation *whatever grounds the disposition to* ϕ . The parallels are knowledge (or ignorance) of other dispositions in nature, as, for example, *whatever causes Parkinson's disease*. Wittgenstein's discussion of the dispositional account of concepts is quite hopeless, therefore, because he considers neither the version that a materialist who was a realist about concepts would hold, nor the version that a mentalistic realist would prefer.

7 Infinite content: Kripke's statement of the problem

Kripke augments Wittgenstein's inadequate attack on the dispositional theory in an important way. Before we can understand Kripke's treatment of dispositions we must look closely at his re-statement of the basic problem. Kripke argues that there is nothing in my mind from which can be read off how I intend to proceed. Hence nothing in my mind determines which concept I am employing:

Let me suppose, for example, that '68+57' is a computation that I have never performed before. Since I have performed—even silently to myself, let alone in my publicly observable behavior—only finitely many computations in the past, such an example

surely exists. In fact, the same finitude guarantees that there is an example exceeding, in both its arguments, all previous computations. I shall assume in what follows that '68 +57' serves for this purpose as well.

I perform the computation, obtaining, of course, the answer '125'. I am confident, perhaps after checking my work, that '125' is the correct answer. It is correct both in the arithmetical sense that 125 is the sum of 68 and 57, and in the metalinguistic sense that 'plus', as I intended to use that word in the past, denoted a function which, when applied to the numbers I called '68' and '57', yields the value 125.

Now suppose I encounter a bizarre sceptic. This sceptic questions my certainty about my answer, in what I just called the 'metalinguistic' sense. Perhaps, he suggests, as I used the term 'plus' in the past, the answer I intended for '68+57' should have been '5'! Of course the sceptic's suggestion is obviously insane. My initial response to such a suggestion might be that the challenger should go back to school and learn to add. Let the challenger, however, continue. After all, he says, if I am now so confident that, as I used the symbol '+', my intention was that '68+57' should turn out to denote 125, this cannot be because I explicitly gave myself instructions that 125 is the result of performing the addition in this particular instance. By hypothesis, I did no such thing. But of course the idea is that, in this new instance, I should apply the very same function or rule that I applied so many times in the past. But who is to say what function this was? In the past I gave myself only a finite number of examples instantiating this function.

All, we have supposed, involved numbers smaller than 57. So perhaps in the past I used 'plus' and '+' to denote a function which I will call 'quus' and symbolize by '+'. It is defined by:

$$\begin{aligned}x + y &= x + y, \text{ if } x, y < 57 \\ &= 5 \text{ otherwise}\end{aligned}$$

Who is to say that this is not the function I previously meant by '+'?¹²

8 Kripke's treatment of the dispositional solution

We can now return to the treatment of dispositions. The hypothesis is that the mental grasp in a realisation of a rule which has an infinite

number of consequences, is a disposition. So what constitutes Kripke's intention to follow the rule for addition, not quaddition, is the presence in his mind or brain of a disposition to follow the rule for adding and the absence of any disposition to follow the rule for quadding. To this Kripke makes two objections.¹³ First, no actual disposition in a person is infinite. Many numbers will be so large that even a competent mathematician will not be able to assimilate them and perform addition with them. Second, many people are disposed to make certain mistakes regularly, although they *intend* to add up correctly. It is the latter point on which Kripke builds. It shows that the subject's intention is a function of what they should do, rather than of what they are in fact disposed to do. If what they *meant* was identical with what they were in fact disposed to do then they mean by 'plus' something which includes all the errors of operation to which they are regularly prone. But someone who is prone to forget to 'carry' when adding certain numbers is not operating with a deviant concept of addition: they mean to add in the normal sense. What is meant must be known in order to judge a disposition: the disposition itself does not constitute the meaning. No doubt someone is susceptible to being interpreted as adding only because their behaviour approximates to making correct additions: in Dennett's jargon, addition is part of the intentional system used to interpret their behaviour. But intentional systems are not reducible to or identical with internal workings or dispositions; rather they are public phenomena.

Kripke's argument only works against a physicalist conception of disposition. If instead it is conceived of as a capacity resting on a direct intellectual apprehension of a meaning or universal, these problems do not apply.¹⁴ But the intellectualist conception means that the identity of the internal state is not a function of the behaviour we in fact produce. Its intellectual content is intrinsic to it, in a way that it cannot be intrinsic to a physical structure, independently of the behaviour it generates. Neither does the infinity of the content of a disposition trouble the mentalist, for direct intellectual apprehension is the apprehension of an intensional entity, and the infinity is identical with the intension.

9 Kripke and the intellectualist solution

Kripke does consider the idea that the apprehension of meaning consists in some sort of direct apprehension. First he considers the

suggestion that this might consist in some introspectible quality, analogous to a headache. This theory is easily dismissed on the grounds that, like the picture of the cube, such a feeling would have no intrinsic meaning. Second he considers the more plausible view that the experience of meaning is not to be assimilated to introspecting a *quality*:

Perhaps we may try to recoup, by arguing that meaning addition by 'plus' is a state even more *sui generis* than we have argued before. Perhaps it is simply a primitive state, not to be assimilated to sensations or headaches or any 'qualitative' states, nor to be assimilated to dispositions, but a state of a unique kind of its own.¹⁵

To this Kripke has two objections. First he says that 'it leaves the nature of this postulated primitive state—the primitive state of 'meaning addition by "plus"—completely mysterious'. On its own, this is no more than a loaded way of saying that it is *sui generis*. But he puts the weight on the second point: 'Even more important is the logical difficulty implicit in Wittgenstein's sceptical argument.'¹⁶ This is the difficulty of the finitude of the mind and the infinity of the content of 'what I mean'. I can see no force in this objection once one has hypothesised that there be a state of intellectually apprehending a universal or a sense. In considering the suggestion that we grasp such senses Kripke says:

ultimately the sceptical problem cannot be evaded, and it arises precisely on the question how the existence in my mind of any mental entity or idea can *constitute* 'grasping' any particular sense rather than another. The idea in my mind is a finite object: can it not be interpreted as determining a quus function, rather than a plus function?¹⁷

I fail to see what, for a dualist, is at stake in the question of whether the idea in my mind is a finite object. It is not infinitely large, because it has no dimensions; it is not infinitely complex, because the idea of addition is quite simple, but these ways of being finite are irrelevant to the infinity of its possible extension. Its extension is infinite, but as we are granting the mind a faculty of grasping intensions, this constitutes no problem. Kripke seems simply to be refusing to consider that grasping an intension may not involve directly grasping its extension, and this is just to reject the intellectualist hypothesis at the outset.

10 The social solution

We have reached something of a clash of intuitions. Neither Wittgenstein nor Kripke appears to have an argument against the intellectualist hypothesis, but modern philosophers might nevertheless prefer to manage without such a hypothesis if they can. Is it at least possible to develop a coherent alternative to accepting as basic the intellectual apprehension of intensional entities? The imagist and the reductively dispositional analyses have already been refuted. Kripke and Wittgenstein have established their case to the point of showing that there are no conceptions of grasping a concept which makes that understanding internal, except for the intellectualist theory. We must, therefore, look for theories which do not locate such understanding internally.

The simplest 'external' theory would be one that said that what concept one meant was a function of what objective similarity in the world was consistently correlated with one's use of a word. The point of such a theory would be to replace the internal apprehension of similarity by the objective common feature of the things to which the word had so far been applied. This, however, is a non-suggestion. So far I have been applying the word 'plus' in a way which maps it into addition-relationships between numbers. But I have also been using it in a way which maps it onto quaddition-relationships. The answer to the question of what similarity or feature in the external world a certain word has so far been applied to is as underdetermined and for the same reasons, as the answer to the question of what I intend or mean.

The argument here resembles Goodman's 'new riddle of induction'.¹⁸ All green objects are, so far, also grue. How do I know that by 'green' I have not so far meant grue? If the argument in the previous section were correct, it would be useless to consult *what I had intended*. The present suggestion is that I mean green because green is the real quality in the world to which I have so far applied the word. But why don't I mean grue? The anti-Goodman answer would be that I don't mean grue because grue is not a real quality in the world, but an artificial one. The same would be said of quaddition. This objection raises difficult, deep and far-reaching issues. Certainly, the fact that similarities can be concocted *ad hoc* does not, of itself, entail that some are not more salient, or more natural, than others. Goodman's own argument that green and blue

are not more basic than *grue* and *bleen*, because green and blue can be defined in terms of *grue* and *bleen*, just as well as vice versa, only shows that salience or basicness cannot be explained simply in terms of primacy of definition. Fortunately, we do not need to become involved in this issue. What someone means by a word is not just a function of how they have used it in the past but also of how they will use it in the future. It would not be plausible to argue that someone who called blue objects '*green*' after the crucial date was using '*green*' to mean green, rather than *grue*. Similarly, someone who made 68 and 57 '*add up to*' 5 was probably always practising quaddition. So if the present hypothesis is to have any plausibility the speaker must not only have previously used the word so as to map it into some salient feature of the world, but must continue to do so. But if they have the propensity to fit the word to the correct feature of the world then presumably they have some sort of internal intention or disposition to do so. It means that they have latched onto that feature as a universal, and so as something with a potentially infinite extension. But it was precisely the supposed impossibility of doing this which refuted the theories which make understanding something internal. '*External*' theories are, therefore, equally forbidden to rely on a correlation between the use of words and features in the environment.

If correctness in the use of a word cannot be explained either by reference to some intention, or to a correlation with how the world is, the only option would appear to be that it consists in using it in a way that others find acceptable: correlation, that is, neither with an internal intention, nor with the world, but with the practices of others. This means that the reason why 68 and 57 are taken to make 125, not 5, is that 125 is the answer which people will accept, and 5 is not. This applies to all judgements concerning what counts as '*the same thing*' or '*doing the same thing*'. Someone who has so far categorised all square things together and who now starts to put round things in that category can be asserted to be wrong only because the practice does not find social acceptance. This does not mean that the agreement in practice is voluntary. Perhaps people do what they do automatically and could not follow if someone tried to do things in certain other ways. But there is no question of any way being right and others wrong, for our concept and practices cannot be thought of as mirroring an objective reality. This is a very radical position, because it commits its protagonist to the

view that the notion of *how the world is*, independently of *how people agree it to be*, has no significance.

11 Walker's refutation of the social solution

The crux of the argument is the conclusion that meaning is the creature of agreement in social practice. Ralph Walker brings two arguments to show that this conclusion must be false.¹⁹ The first is that it will not do to treat agreement as a bed-rock fact that requires no explanation: why do we agree that certain things count as 'the same', particularly given that we seem to light upon genuinely projectable categories, and not ones that have proved—so far—to be grue-like? In Walker's words:

Wittgenstein's whole account of language and of truth depends upon the continued agreement among speakers as to which things are to be classified as red, square, etc. To say that this agreement is inexplicable is preposterous: it is to make the whole of language and coherent thought depend on a continued coincidence, the accident that our judgments of similarity go on agreeing as new instances are encountered. In saying it is preposterous I am not saying it is self-contradictory; there is indeed no contradiction in supposing that so vast and inexplicable a coincidence has been occurring and making language possible for us. But a philosophical theory can be wild without being self-contradictory.²⁰

Walker's second objection does not rest on an intuition. It is more strictly formal. If facts are created by agreement then *the fact that something is agreed* to be a similarity will itself be created by agreement, and so will the fact of this latter agreement. The problem is that the theory appears to take the *social* fact of agreement in practice as somehow given and *from that* to construct the facts—the objective similarities—in the world. But *what is to count as the practice of agreeing* or what is *sameness of practice* and *that we have an agreement in practice in a particular case* would have to be sustained in just the same way as any other kinds of sameness—that is, by social practice and agreement; and what constitutes these would be further agreement, and so on. It is obvious that we have a regress here: the only question is whether it is vicious. I agree with Walker that it is. If what constitutes agreement in practice is agreement in

practice, and so on, unless there actually is, as a brute fact somewhere in the regress, agreement in practice, then none of the other levels can get made or constituted. It is not possible to do without brute facts that precede our activity of apprehending the world. This requirement is not satisfied by 'internal realism', for the internal facts are those that are constructed by agreement: what Walker shows is that we need facts presupposed by the analytical account itself, and they are, therefore, externally real.²¹

The conclusion is that what constitutes 'the same' cannot essentially be a product of what people collectively agree to accept. That it was this, however, was the fundamental step in proving that there could be no determinate nature to the logically private. This version of the anti-private language argument is, therefore, no more successful than the traditional one.

12 Not an argument but an alternative account of language

At the outset of this chapter I suggested that the anti-private language argument need not, despite its title, be seen so much as an argument as a proposal for a radical understanding of language and its relation to mind. The motivation for adopting this new model would be that it avoided all the problems associated with the picture of the mind as a private arena. This Cartesian picture has been associated with sceptical problems and with arguments for dualism that disturb the materialist consensus. The failure of the arguments against private languages does not show that its conclusions cannot be adopted on these more general grounds. What I mean by 'its conclusions' is not merely the assertion that there cannot be logically private objects—that alone without argument would be too naked an assertion—but some overall conception of how language works that would rule out the logically private. This will only be impossible if the objections I have presented to the particular forms of the argument include decisive considerations against the truth of its conclusions, rather than merely against these particular ways of trying to reach the conclusions. This, however, is the case. The Kripkean conception of language has been shown to be incoherent, so that cannot be the model of language which excludes the private that we could adopt. And the picture of language as requiring public justification—resting on an undefended public verification—was shown to be inconsistent with experience's role in informing us about the world. It looks, therefore, as if there is not merely no *proof* of the

Wittgensteinian conclusion, but that there is no coherent account of language of which it is a part: no known view of language which excludes reference to the private is coherent.

13 Memory, ‘seeming’ and privacy: an appendix to the traditional argument

In *Philosophical Investigations* 258 Wittgenstein focuses on the problem of recognising a private sensation when it recurs, and the difference between remembering it correctly and only seeming to do so. Because of the nature of the example in the *locus classicus*, it came to be assumed—implicitly at least—that there was an important connection between the insufficiency of unsupported ‘seemings’ and memory. We all feel guilty about our absent-mindedness, and, when this is linked to the shadowy ontological status of the past, it lends plausibility to the idea that private memory ‘seemings’ lack facticity. It is important to realise, however, that memory has nothing essential to do with the argument. The problem would have been exactly the same for the private diarist if they had been trying to decide whether two *simultaneous* sensations—say, one in their right hand and one in their left—were sensations of the same kind. The situation is the same because it relies equally on how it uncheckably seems to them. The conclusion would have equally to be that there could be no fact of the matter. Because there is not the distance in time that memory imports, it is clearer in this case that the ‘seeming’ in question is the phenomenological one, rather than that of mere judgement; and in this case the idea that one cannot tell seems manifestly absurd.

It is not possible to allow that judgements about simultaneous sensations are acceptable whilst disallowing ones that are diachronic. Someone might try to operate this distinction, on the grounds that simultaneous recognition alone could not establish meaning, for language must operate diachronically. But this is not the point. Once it is conceded that there is something for the simultaneous judgements to be right about, then there are the appropriate private objects and, therefore, something about which future memory claims can be right or wrong. Furthermore, there is no clear distinction between present recognition and memory: if I judge that two simultaneous sensations remain the same for two seconds, is that a memory judgement or a direct one?

Perception

The simultaneous case seems to me to be the real intuitive test for the credibility of Wittgenstein's conclusion. Only a really powerful argument could overcome the intuition that we could *obviously* tell under these circumstances. We have seen that there are no overwhelming arguments for this bizarre conclusion.

CHAPTER V

Contemporary Physicalist Theories of Perception

1 Experience and reductionism

We have followed the problem of perception through the arguments for sense-data and how they have been resisted, and through the Wittgensteinian attack on the empiricist conception of experience. If one is impressed by these attacks on traditional empiricism then one will be left with a choice between some form of direct realism, or some other theory that avoids the Cartesian theatre of experience. How modern direct realist theories fare will be the subject for Chapters VI and VII, but the paradigmatically contemporary flight from sense-data takes the form of physicalist theories of perception.

Physicalism is a theory about what there is, and the philosophy of perception concerns a certain way of knowing about the world: physicalism is, therefore, an ontological theory and the philosophy of perception is a branch of epistemology. It might, therefore, be thought that the two were independent. The matter is not, however, so simple.

Theories in the philosophy of perception that belong to the sense-datum family postulate private objects of awareness that possess such sensible qualities as colour, shape, sound, felt texture and the like. Such objects are not, however, to be found in the brain and, if there were little green patches, smells and noises observable in the brain this would still not tell us what it is to be *aware* of them. It did not take materialists long to realise, therefore, that phenomenal objects are part of the Cartesian conception of consciousness as a non-physical inner arena, and that they are consequently something that they must reject. Belief in phenomenal objects is, therefore, regarded by physicalists as a howler. U.T.Place, in a classic paper, dubs such a belief ‘the phenomenological fallacy’; this is the (supposed) fallacy of treating the

contents of experience as entities which really possess the qualities that they seem to possess.¹

Because they reject the Cartesian arena, physicalists have tended to welcome the conclusion of Wittgenstein's argument, even when they have not been convinced by the argument itself. That is, even if not convinced that there is anything *conceptually* wrong with the notion of the logically private, they are committed to the view that there is no such realm. Reductive materialism cannot allow the existence of a private arena of sense-contents, or a special subjective dimension: everything must be in the public world of physics.

It might seem that rejection of the sense-datum theory would push the physicalist back to naive realism. But naive realism is hardly more compatible with physicalism. The naive realist believes in direct consciousness of the external world, and takes this consciousness to be something basic. But consciousness is a mentalistic notion. The physicalist's approved set of relations are all spatiotemporal and/or causal, but the naive realist does not believe that the perceptual relation is *simply* a matter of a physical object outside an observer causing some sort of state internal to them: the *experience* of the object may depend intimately, or supervene, on such a causal process, but it is something more than such a process. Both naive realism and the sense-datum theory affirm that experience possesses a special felt nature which is not captured by any physical description: the difference between them concerns only the object of this *sui generis* feeling—one holding it to be private and the other public—and this is, in a sense, a domestic dispute amongst dualists. It is the *sui generis* conception of experiential feel that the physicalist must reject, for, whatever the object of such feeling, *how it is with the subject* is something private to them and, hence, not capturable in physical science.

In order to meet this requirement, some sort of physicalist analysis of experience is required. The crudest such attempt is behaviourism, according to which perception is simply the physical stimulation of a sense-organ in a way that gives rise to a disposition to respond in an 'appropriate' way. Such an account seems too bare—a behavioural disposition is too abstract to do the job of the inner 'raw feel' of experience.

This perceived inadequacy of a simple behaviourist approach to perception led to two developments. One was to central state materialism. According to this theory, mental states are brain states.

The 'raw feel' of experience can, therefore, be identified with something current and concrete, rather than something abstract. Of course, although, when we are aware of our 'raw feels' we are in fact aware of brain states, we do not experience them *as* brain states: but neither, for the physicalist, do we experience them as possessing explicitly mental properties—that is, properties different from those found in the rest of the physical world. We experience them *topic-neutrally*, as things possessing a certain role, without knowing what the things possessing this role actually are in their intrinsic nature. Science comes along to tell us that it is—probably—brain events that fill this role.²

The topic-neutral approach to experience will be discussed on pp. 136–8, but it is the other post-behaviourist development which will occupy the first half of the chapter. This is what might be called the *purely cognitivist* theory of perception.

There are two general categories of mentality which pose problems for physicalism, namely the sensory and the intellectual. It seems generally to have been agreed that the intellectual was more susceptible to a roughly behaviourist analysis than the sensory. It is less obvious that treating belief, for example, as some sort of dispositional state, leaves something out, than that a similar treatment of sensation omits something: in sensation one seems directly to confront non-dispositional content in a way that one does not so plainly do in believing. This raises the thought that if sense-experience could be analysed in intellectual or cognitive terms then it might prove more susceptible to a plausible reduction.³

This is the path down which many modern physicalists, such as Armstrong, Dretske and Pitcher, go, first analysing perception as no more than the acquisition of beliefs or information through the senses and then analysing belief and the possession of information in causal or dispositional terms. Armstrong states the position as follows:

In the first stage, it is argued that an account of perception can be given in terms of the acquiring of beliefs about the physical world... Such an account, however, must take as primitive the psychological concept of belief. It is therefore incomplete from the point of view of the attempt to give an account of mental states simply as states of the person apt for the production of certain physical behaviour, or states apt for being brought about

by certain physical objects or situations. The second stage of the argument tries to show that the acquirings of belief involved in perception are susceptible of this sort of analysis.⁴

Nowadays, the reduction of perception to cognition without remainder—what is generally called a purely cognitive theory of perception—is only adopted as the first stage of this kind of physicalist reduction. But this need not in principle be so. It is possible to hold a purely intellectual theory of perception without wanting to reduce the intellectual. Such a theory was held by certain rationalists and absolute idealists. These other theories will also be covered by the first part of the discussion that follows.

2 Pure cognitivism

I have explained what I mean by the expression ‘a purely cognitivist theory of perception’ and that such a theory can be considered either in its own right or as a preliminary step to a materialist theory of perception. That there must be something wrong with the assimilation of the experiential to the cognitive is shown by considering what I call *the principle of minimal empiricism* (henceforward PME).⁵

PME says that, in the end, our grasp on concepts depends on the fact that we experience instances of at least some of them. It need not be assumed that there are any particular concepts that must be cashed in this direct way. The principle is well expressed by Feigl:

Our world, being what it is, can, of course, be known by description in any of its parts or aspects, but only on the basis of a foothold somewhere in acquaintance. This, it seems to me, is one of the cornerstones of any empiricist epistemology old or new. But the new empiricism of recent times has come to recognise that it matters little just which areas of acquaintance are available or actually utilised... the congenitally blind-deaf person... could in principle construct and confirm a complete system of natural science.⁶

A stronger form of empiricism may be justifiable, but this is all that is needed for the principle of *minimal* empiricism. To deny this principle would be to say that our understanding of concepts could rely entirely for its content on their interrelation or inter-definition with each other, for this is all that is left if we rule out our grasp on their relation to the

world—such a grasp being available only through sense-experience. And to rely on their interrelatedness is to deny that there is any difference between concepts that constitute a purely formal calculus and an interpreted calculus. I take it that it is an absurdity to deny this distinction.

The question is, therefore, how perception performs this task of giving an interpretation to our concepts. There are two possible answers to this. The more natural one is to say that in perception one confronts instances of concepts in a way that one does not do in other kinds of conceptual activity: perception is *presentational* in a way that pure thought is not. It is because perception presents features of the world that it is possible to go away and think about them: presentation tells us what certain empirical features really are like. By seeming to see something red I can learn what red is, or is like; it would not be possible to acquire such a concept by pure thought. The other, and less natural, explanation of how perception interprets our concepts is the behaviouristic one; it is from the way that our concepts are integrated with our dispositions—in particular our capacities for discriminatory dispositions—that they are interpreted empirically.

It is difficult to see how a cognitivist could adopt the *presentational* explanation of the PME, because it seems to rest on contrasting the way in which objects figure in experience with the way they figure in purely cognitive activity, whereas cognitivism seeks to assimilate the two essentially, letting any difference rest on such extrinsic facts as causal origin: for the cognitivist, perception is nothing but the acquisition of a belief, but one that happens to come through the senses. Some cognitivists have tried to explain the phenomenological difference between perception and more obviously purely conceptual activities by saying that it is the sheer quantity and density of information that give perception its ‘fuller’ feeling. Perception is rather like an old-fashioned newspaper photograph, made up from a mass of punctiform information relating to every point in the sensory field. Given that it is the essence of cognitivism to treat perception as no more than a kind of involuntary thought, the cognitivist would seem to have no other resources to use to explain the apparent phenomenal difference between perception and thought, except by something like a difference in quantity.

This theory—which we might call ‘the density theory’—will not do. If we take red to be one of those concepts that a particular person

understands through its connection with perceptual experience, and then say that perception of red differs from pure thought of red by the fact that the perception consists of a particularly dense collection of thoughts with the content (non-linguistically expressed, of course) 'this is red', then we will be moving in a circle. On the one hand, the pieces of information that make up the dense perceptual bundle will each contain the concept 'red', whilst, on the other, that concept only gets its content from the perceptual situation. Conceptual activity is, thus, both more primitive than perception, because it is that out of which perception is constructed, and derivative from it, because it relies on perception to give it its content. It can be shown more formally that this is viciously circular. The first premise in the argument is just the PME:

- 1 It is a necessary condition for anyone's having the ability to comprehend cognitive content that they have a grasp on some sensory content or other.

The second premise is the density theory of the nature of sensory content:

- 2 The contents of sensory acts are constructions from those of cognitive acts: e.g. seeming to see something red is to be analysed as acquiring beliefs or information, through certain channels, about the presence of something red.

The third premise is a general principle about the relation of constructions or analysanda to their elements or analysantia:

- 3 If *As* are constructions from *Bs* then it cannot be necessary that someone comprehend *As* as a requirement for comprehending *Bs*.

Therefore

- 4 It cannot be necessary that someone comprehend the contents of sensory acts as a requirement for comprehending the contents of cognitive acts. ((2), (3), by instantiation.)

(1) and (4) are contradictories, so one of the premises must go. The PME has already been defended, and (3) seems obviously sound—provided that the necessity involved be not just empirical or epistemological. But the necessities in all the premises are broadly conceptual. The only premise left vulnerable is (2), the density theory.

What this shows is that perceptual experience must have a core which is pre-conceptual.

3 The perceptual core

The options

It might seem that there are only two ways in which one could understand the idea of a pre-conceptual core to experience. The most obvious is the traditional one, which affirms the existence of a phenomenal realm in some non-reductive sense. This is just what the physicalist is trying to avoid. The other is to treat the preconceptual as consisting in very basic behavioural reactions. That is to say, the interpretation of our concepts through perception happens because our basic conceptual activity is rooted in some fundamental behavioural responses—basic discriminations of features in the world. Only when these responses reach a certain level of complexity do they count as conceptual. Such a theory will be part of a behaviourist-cum-functional account of concept possession. This is the position that is represented by functionalist as well as behaviouristic theories of mind.

I believe that it is correct that the behaviouristic strategy is the only one available to the physicalist. There are, however, various ways of talking about perception current among materialists which could give the impression that there are other options available. I say above that many physicalists fear that the behavioural or functional framework is too abstract to capture experience alone, and would prefer to give—perhaps within that framework—a physicalistically respectable sense to the phenomenal. Paul and Patricia Churchland give a physicalist interpretation to the notion of *qualia*, which, if it worked, could be thought of as doing just this. Others seem to think that there is a concept of *pre-conceptual information* which, perhaps aided by the notion of *analogue representation*, can do the job. I shall discuss the Churchlands' theory first.

The Churchlands' qualia

The Churchlands try to save the phenomenal for physicalism by developing a physicalistically respectable notion of *qualia*. If this could be done successfully, then a pre-conceptual kind of representation—a

form of knowledge by acquaintance—should be available to explain the PME. The Churchlands define *qualia*, uncontroversially enough, as follows:

‘Qualia’ is a philosophers’ term of art denoting those intrinsic or monadic properties of our sensations discriminated in introspection. The quale of a sensation is typically contrasted with its causal, relational, or functional features.⁷

Because of the connection with experience, the natural way of understanding ‘discriminated in introspection’ is to take it as meaning that in introspection one is aware of, such that one can recognise, the qualitative nature of the *quale*. The point of *qualia* is that they constitute the qualitative content of experience: *qualia* are qualitative entities and being introspectively aware of them would seem to imply being aware of their qualitative natures. It is by their presenting empirical qualities that *qualia* make possible the empirical interpretation of our language. The Churchlands, however, are using ‘discriminate’ in a purely extensional sense, in which one can discriminate something without being in a position in principle to recognise what it is that one is discriminating. In their view, the intrinsic properties of the sensations turn out to be such things as

the spiking frequency of the signal in some neural pathway, the voltage across a polarized membrane, the temporary deficit of some neurochemical, or the binary configuration of a set of direct-current impulses.⁸

Given that such different physical realisations are not experientially discernible as such, and yet they are supposedly the *qualia* themselves, it follows that one cannot tell by experience what *quale* one is experiencing. In the Churchlands’ own words, they are only ‘*opaquely* discriminated’.⁹ If two different physical states played the role of the same kind of sensation in the same brain they would be totally different *qualia* (because their monadic properties are different) but no difference would be available to introspection, for the physical properties are not themselves discernible.

This is a travesty of the idea of a *quale*. If the term *qualia* has any use at all it is to designate *experientable* differences; significantly different *qualia* can, in principle, be recognised as different. The point is not purely verbal. If *qualia* do not constitute the available content of experience they make no contribution to explaining the

difference between experiences and other so-called representations. According to the Churchlands, purely conceptual activity is also identical with neural performances of one kind or another. *Qualia* are meant to distinguish the perceptual from such less vivid cognitive states. They will do this only if they constitute the content of experience in some especially direct way. Nothing the Churchlands say explains how they do this.

A further confusion in the Churchlands' understanding of experience is relevant. They are emphatic about the plasticity of experience, by which they mean that the content of experience is relative to the framework by which it is described and interpreted; it is not 'given'. If we were taught to report on our experiences using such expressions as 'the spiking frequency in pathway x is n' then we could be said to be directly aware of such processes as the spikings. However, they also say—and obviously correctly—'our mechanisms of sensory discrimination are of insufficient resolution to reveal on their own' such detailed facts as what goes on in the brain.¹⁰ Far from showing either that experience is relative to whatever theory one happens to hold, or that it is of whatever micro-process happens in fact to be its cause, what this suggests is that experience is most faithfully expressed by a description or conceptualisation that has just the same grain of accuracy as the sensory or introspective mechanism in which it is embodied or 'realised'. A description at this level is bound to be topic-neutral about what the actual process is. The only plausible candidates for this are the functional ones used by Smart and Armstrong, of the general form 'something is going on of the sort that goes on when there is a certain stimulus', or 'that is apt to cause a certain response'. One is not now discriminating a monadic property opaquely, but discriminating functional properties explicitly, and the theory ceases to be relevantly different from functionalism in general, and will be subject to whatever objections we can bring against such theories.

Analogue representation

Dretske says that he wishes

to develop the idea that the difference between our perceptual experience, the experience that constitutes our seeing and hearing things, and the knowledge (or belief) that is normally

consequent on that experience is, fundamentally, a coding difference.¹¹

The difference in coding is the difference between *analogue* and *digital* representations. An analogue representation is distinguished by two closely connected features. It represents in a continuum, rather than discretely, as a normal watch or speedometer does in contrast to a digital one: and it carries a mass of information, rather than a selected piece, in the way that a photograph does. Digitalising information is abstracting a specific piece from the continuous analogue mass, like making a verbal report of some feature in a photograph. Experiences are analogue representations.

All this is doubtless true. A sense-datum, like a photograph, is an analogue representation. What is relevantly important is what the reductive physicalist can make of it. For present purposes, the main issue is whether it offers some alternative to a functionalist interpretation of experience. Physicalistically, an analogue representation would be a brain structure that responded more or less continuously to some continuum of external properties. There is no more problem with this, in principle, than with a thermometer reacting to heat or a speedometer to speed. The theory loses its interest in the present context, however, when we consider what makes such a structure *count as* a representation of some external property. This can be nothing other than its functional role in being standardly brought about by such a property and playing a role in modifying our reaction to the property. The doctrine of analogue representation, in a physicalist context, is firmly embedded in functionalism: it constitutes an account of how a certain kind of function might be realised, but it does not augment or modify the pure functionalism of the conceptual account.¹²

The behaviouristic approach

It has proved correct to say that the behaviourist-cum-functionalist approach to experience is the only one open to the reductive physicalist.

A behaviourist-cum-functionalist (henceforward, just 'behaviouristic') approach to the PME faces many problems. It faces, *a fortiori*, all the general objections to such theories. The most famous of these in connection with their ability to accommodate perceptual

experience is the '*qualia* problem'—that is, the problem raised by the fact that such theories seem to omit the qualitative content, or 'raw feel' of experience. More directly concerned with the PME is the following difficulty.

The PME both has a general appeal, based on the fact that without it nothing would count as an empirical interpretation of our concepts, and a first-personal appeal. When I, as a subject, think of my understanding of empirical notions, the fact that certain concepts can be matched with certain kinds of experience plays a vital role in my own sense of what I mean by my language. This sense comes from putting together word and object of experience, not from putting together word and nexus of behavioural responses. If I want to understand 'red', for example, I fix my mind on an appropriate phenomenal object of experience; I do not try to concentrate on how I am disposed to react to it. Someone else, however, who was teaching me the idea or trying to test my knowledge of it, would concentrate on my reactions. The behavioural understanding of empirical interpretation is essentially third-personal.

It is a common intuitive objection to behaviouristic theories that they are viciously third-personal; that is, that they present a third-personal perspective on something which is essentially first-personal, namely, the viewpoint of the conscious subject. As with many fundamental disagreements—especially in the philosophy of mind—there is great difficulty in turning what appears to be a clear intuition into a demonstrative argument—that is, into an argument that does not rest on a premise which is more or less equivalent to, and just as contentious as, the initial intuition. I am going to try to turn the intuition that behaviouristic theories are viciously third-personal into an argument by showing that such theories cannot be applied to the first-personal perspective without a vicious circularity.

The problem can be brought out by considering the sort of situation that is the classic context for behaviourist interpretation: animal behaviour. It seems fairly natural to interpret the understanding of the world possessed by a rat in a maze as being no more than its ability to react to the barriers it encounters: one need not think of anything else inside the rat constituting its understanding other than those dispositions. But we can only make sense of the rat's understanding in these reduced terms by interpreting those dispositions against a background picture of how the world around the rat really is: we have a non-reductive grasp on how the world really is and how the rat actually

behaves, and against this background we can see the rat's grasp on the world as no more than its dispositions to move in certain ways. It would make no sense to try to understand the rat's dispositions in abstraction from a general conception of the physical world, for the notion of a behavioural disposition is locked into that conception. It involves concepts such as body, motion and spatial location, which are constitutive of our conception of the physical. We can see the rat's cognition as merely dispositional because we take our own cognition naively.

The situation of a rat is no different from the way we stand with regard to other minds in general; we employ our non-reductive grasp on the world to set up the framework in which we can make the reductive interpretation of others. But we are no different from others, and if the reductive theory were true we ought to be able to apply it to ourselves. Yet, as the case of the rat clearly shows, seeing cognitive grasp in terms of dispositions necessarily presupposes a direct grasp on the world, both because the naive picture of the world forms the backdrop against which dispositions are understood, and because the concept of disposition can only be understood in terms which presuppose a grasp on basic physical concepts. One could not reduce one's own apprehension to dispositions without a regress, for 'disposition' is a concept which only makes sense as part of a general conception of the physical world. So I can never make the reductive interpretation of my own understanding, paraphrasing my conception of the physical in dispositional terms.

My attempt to make this intuition more rigorous involves four principles.¹³ One of these has been given already. It is the *non-basicness of the concept of disposition*. The other three are (1) the doctrine that *thought is the vehicle of understanding*, (TVU), (2) the *entrapment of content* (EC), and (3) the *eliminability of reductive analysanda* (ER).

1 (TVU) One understands the world by means of one's thinking. There is a difference, that is, between the objects of cognition and of concepts, on the one hand, and the acts of cognising and conceiving, on the other; and it is by means of the latter that we grasp the former. Thought may not be essentially verbal, but there is an analogy between thought and language. One can use language to express thoughts—one can really talk about things—only if one understands the language. Similarly, one can think about things only if one understands one's thoughts (generally and paradigmatically, at least). So there is a parallel

between the senses in which language and thought are vehicles for understanding, namely that in both cases the vehicle must be understood in order for the object or information to be grasped through it. Even though thought is essentially transparent in a way that language is not (the concept 'red' could not have been another concept but the word 'red' could have had another meaning, so you cannot fail to understand thought in the way you can fail to understand language) it is still true that thought is the psychological vehicle by which its object is conveyed; there is still a difference between a thought *qua* psychic episode and its object—even if the episode is thought of as consisting in the object intentionally realised. And it is because we can understand the psychological episodes which are our thoughts that we are able to understand—form a conception of how it is with—the world.

2 The point of EC is that a psychological state includes its conceptual or propositional content, so that any account of the nature of such states will comprehend their content. What is at stake here is whether a reductive analysis of understanding also affects our grasp on what is understood. Physicalists who reject this intuition object that when they present an analysis of what it is to understand some concept or proposition *c*, it is the *understanding* they are analysing *not c*, its content; so, they argue, saying that our understanding of the concept *body*, for example, is simply a disposition has no consequences for what we understand by, or in, that concept. By contrast, EC claims that, for cognitive states, *the essence of which is their content*, this distinction cannot be made; a reductive account of understanding *c* ought to be such as to entail that the cognitive act does indeed have that content. If a certain state contains information of some kind, then it ought to follow from a reductive description of that state that it contains that information—if the reduction is correct. (It might seem that this simply denies an externalist account of content. We shall see on p. 136–50 that externalism is a multiply-ambiguous idea, but it is irrelevant in the present context. Whatever mental content might be external too, it is not external to a full specification of the dispositional content in terms of which the reductive account is given, for such a disposition is sufficient to determine the specific content of that state.)

3 ER states that, if a reductive analysis is correct, then it ought to be possible to operate entirely in terms of the reductive account. For example, if knowledge is justified true belief then it ought to be

possible to drop the notion of knowledge in favour of justified true belief. This is so because reductionism is like eliminativism in that it is a condition of its adequacy that one be able to conduct all one's business, in principle, in the preferred terms. The reductionist looks at their analysans and sees that it entails that which they are reducing: the eliminativist considers their preferred theory and sees that it does everything legitimate that was done by the eliminated theory, but that it does not entail it. They are equal in finding their preferred discourse omniscient.

These three principles come together with the fact that the concept of disposition presupposes more fundamental physical concepts to generate an argument against dispositional theories of cognising or understanding. Put informally, the argument is that if our thoughts and cognising are understood reductively, then they no longer have the kind of content that could make the world available to us; rather the content so conceived presupposes a grasp on the world (that is, a conception of the world; its correctness is irrelevant). The cognising, *qua* psychological state, could no longer be thought of as the vehicle by which a grasp on the world comes, rather cognising would presuppose such a grasp. The moral is that, in any adequate account of cognition, its object must be thought of as present in the psychological state in some manner which is not reductive.

We can now put together the following argument. The idea that thought is the vehicle for understanding can be expressed in the following way:

- 1 I can understand the subject-matter of my understanding—i.e. the world—only by means of my understanding, and, trivially, this presupposes I am able to understand the content of my understanding.

Therefore, *a fortiori*,

- 2 My being able to understand the subject-matter of my understanding—i.e. the world—is by a means which *presupposes* I am able to understand the content of my understanding.

We now bring ER and EC into play:

- 3 If a reductive account is true then everything relating to its topic can be considered in the reductive terms without loss—for that is what there really is.

- 4 The topic of a reductive analysis of cognitive states includes their cognitive or informational content; that is, that a state has the content it does must follow from the analysis provided of it.

These principles combine to give

- 5 If a reductive account of understanding is correct then anything relating to its content can be considered in reductive terms without loss.

So (2) is a specification of TVU and (5) follows from combining ER and EC. Together they give

- 6 If a reductive account of understanding is correct then my being able to understand the subject-matter of my understanding—i.e. the world—is by a means which *presupposes* I am able to understand the content of my understanding in the terms in which the reductionist analysis represents it—that is, as a disposition to certain kinds of physical behaviour.

But this brings trouble from the fact that ‘disposition’ is non-basic, for this idea is expressed in (7):

- 7 Because the concept of disposition presupposes a whole battery of physical concepts, I can understand the states which are, according to the reductive account, my understandings of the world, only if I already understand the subject-matter of my understanding (that is, the world itself).

(7) says that the reductive understanding of content—in terms of having dispositions to behave—presupposes understanding the physical world—that is, presupposes a conception of the world. (6) says that, if reductionism is true, we get our conception of the world by means of the reductive understanding of content. So we get our understanding of the nature of the world by means of something which presupposes we already possess it. This is obviously a contradiction, so (7) and the consequent of (6) are inconsistent. Given that (7) is a simple analytic truth, we can, therefore, work a *modus tollens* on (6) and conclude

- 8 The reductionist account of understanding is false.

I think that the only possible objection to this argument is to challenge the generality of (4). That statement of the *eliminability* of

reductanda claims that reductanda are always replaceable by the reductans. One would not expect such a replacement to be possible in opaque contexts: for example, one would not expect it to follow from the fact that Smith believed that he understood the predicate calculus that Smith believed that he was disposed to behave...(now fill in the behaviour appropriate to understanding the predicate calculus). This is relevant because it looks as if the context in the argument might be opaque. I must understand the content of my understanding, but must I understand it under any particular description—in particular, must I understand it under the description that expresses the reductive analysis?

Although this objection is plausible, it is mistaken. If a reduction is correct then there can be no conceptual necessity in operating in terms of the analysandum. The analysans presents what is actually going on, such that it ought to be possible that a creature operate *ab initia* in those terms: there ought to be a creature, that is, that never thought in the terms of the analysandum. The situation is that the objector wants to give priority to the folk psychological description of content, even though, on their own theory, that is not the most fundamental description of what is going on. Furthermore, opacity cannot be relevant. What is at stake is whether the state contains certain information, and a more fundamental characterisation, of the sort a reductive account is supposed to give, cannot result in loss of information.

The position of the objector is analogous to, or, perhaps, even essentially the same as, Dennett's in 'True believers'. Dennett argues that even an omniscient observer, who was able to predict the behaviour of humans by predicting the behaviour of the individual atoms that make them up, would need folk psychology. They would need it if they wished to understand the utterances of humans when they talked to them, and, more fundamentally, they would need it to understand what they themselves were doing. So the folk psychological level of description is ineliminable, though it carries no fundamental ontological clout.

The problem with Dennett's position is that there can be no explanation of why we must adopt the folk psychological perspective. If we are all just clouds of atoms, why are we obliged to see ourselves in this particular ontologically non-basic way? It is true that we cannot see ourselves as *people* or *understand* our *actions* unless we adopt this perspective, but why see ourselves in these ways? An

eliminativist would argue that it is just conceptual conservatism. But if one rejects the idea that we just happen to be hooked on this way of seeing ourselves and agree that the applicability of these categories is truly fundamental, then there is the problem of explaining why it should be compulsory to adopt this level of discourse. A reductionist believes that statements on this level can be true, because they are reducible; this gives them some advantage over the eliminativist, who thinks it all false. But this fact does not explain why, amongst all possible non-basic levels of discourse, this one should be unavoidable, rather than merely available if required. It is possible to argue that the question 'why should we see ourselves as persons?' answers itself, because the use of 'we' already presupposes the personal perspective. But this misses the point. The behaviour of the physical structures that we call 'people' cannot be understood in a way that seems complete or remotely adequate without the personal perspective. Even if we were able to do all the predicting that physical omniscience would make possible, it would be impossible to restrict one's understanding of oneself to the physical terms. The Cartesian certainty that I *think* is absolute, not relative to adopting one possible but inessential level of discourse. Our existence on the personal level is a fundamental, not a pragmatic, fact. There is no way it can be thought of as a function of a certain way of thinking or conceptualising: it is a basic fact, which would not be basic if the physicalist ontology were correct.

That we are people, and that folk psychological concepts apply to us, would have to be a bed-rock and irreducible fact about us, and should not be explained by the fact that the intentional level is one of many real, but ontologically non-basic, levels of description.

There is a reservation of a different sort that might be felt about this argument. It is directed against an across-the-board behaviouristic analysis of cognition. What we were discussing was the behaviouristic analysis of experience. The falsehood of a general behaviouristic theory does not entail that such a theory is not correct for sensation and perception. Reflection on the nature of that argument, however, reveals that it refutes the behaviouristic treatment of sensation in particular. That argument showed that knowledge of the external world cannot be reduced to dispositions because the very idea of a disposition functions only in the context of an unreduced grasp on the physical world. But there is nowhere else that we might get our conception of the physical world from, other than perception.

It is the content specifically of perception that must be taken non-reductively, if the contrast required by our concept of disposition is to be maintained.

4 Physicalism and externalism

The rest of this chapter is concerned with quite a different tack in the assault against physicalist theories of perceptual content. It is directed against mind-brain identity versions of reductive physicalism, by which I mean any physicalist theory that tries to retain the Cartesian common-sense intuition that experience is internal to the subject. It is simply stated: physicalism of this kind requires an externalist approach to perceptual content and such an externalist account of perception is demonstrably false; therefore, by *modus tollens*, physicalism of this kind is false too.

The validity of the argument is patent, so we need attend only to the premises. The first is that physicalism requires an externalist approach to perception.

An internalist theory of perceptual content states that content consists in some kind of internal object of the experience. An internalist is someone who thinks that content is, in general, minddependent. In an intuitively clear sense, the contents of an hallucination or an after-image are internal, because the content of these experiences, however analysed, have their *esse* as *percipi*: we can say that a perceptual internalist treats all content, in this respect, like after-images and hallucinations. The traditional sense-datum or Lockean ideas are paradigms of such contents, but internal contents have not always been conceived in quite these act-object terms. At least some adverbial understandings of content are—in intention at least—internalist and so are some interpretations of the intentional object theory. The kind of physicalism which is relevant to our discussion of perception is precisely concerned to reject this sort of private content. Furthermore, it is generally conceded by reductive physicalists that a state of the brain—or of any other physical object—cannot be intrinsically *about* anything, for *aboutness*—or intentionality—is not an intrinsic physical property of anything; so there can be no internal objects for a physicalist.

By contrast, externalist theories of content for a given mental state say that the content is constituted, in part at least, by something external to the mind: the mental state, in the jargon, is object-dependent. The

simplest externalist account of perception is naive realism. Here the external object (perhaps one should say its perceptible or perceived features) constitutes the content of the experience. In so far as ordinary naive realism takes perceptual consciousness as a primitive, not to be given some reductive analysis, this theory is not open to the kind of physicalism we are discussing.

The physicalist can neither allow the content to be internal, nor allow the sort of externalism that preserves a mentalistic notion of consciousness. The relation between the content and the internal state that is deemed to be the awareness or perception of it has to be an acceptably physical one. This means that, in effect, it has to be a causal relation; for spatial, temporal and causal relations are the only respectable physicalist candidates, and space and time alone are not enough. The only real candidate for such an account is something on the following lines:

X is an experience as of F if x is that state of the subject
standardly caused, through the appropriate sense organ, by Fs.

As this is the only theory open to the reductive physicalist, it is clear that the first premise of our argument—that they must be an externalist about perceptual content—is correct. It is also worth noting that this formula is essentially the same as the topic-neutral characterisation of experience which is employed by central state materialists. Our conception of experience is only of it as ‘that state, *whatever it may be*, that is standardly caused...’. The topic-neutralist approach to experience is externalist: lacking any conception of experience’s intrinsic nature, we are forced to characterise its content indirectly, via some causal relation.

This kind of formula is popular amongst physicalists, particularly those operating within the framework of cognitive science. It is, indeed, just the application of a *causal theory of semantics* to perception. For example, Burge, following Marr’s theory of vision, says that

the information carried by representations—their intentional content—is individuated in terms of the specific distal causal antecedents in the physical world that the information is about and that the representations normally apply to.¹⁴

It is necessary to get clear on one potential area of misunderstanding, however. These accounts are presented as accounts of what it is for

something to be a representation or to have *representational content*. In other words, they concern what it is for some kind of inner state genuinely to represent some feature of the external world. Now the formulae we have been considering could be accepted in that role by a sense-datum theorist who was a representative realist. Someone who believed they saw red the way you saw blue could agree that their 'red' sense-datum and your 'blue' one represented the same external feature. (Whether that feature is to be deemed red or blue will depend on the standard public use of the language.) So a causal theory of representation could be invoked to bridge the gap between the phenomenal and the physical worlds. But this is not their role within the physicalist enterprise. The physicalist is particularly embarrassed by phenomenal properties, and the externalism is needed to explain phenomenal content and not just representational content. For the traditional representationalist, the inner, phenomenal world of the 'manifest image' is projected onto the outer world: for the reductive physicalist, there is no inner world to project. On the computational model, representations are intrinsically only 'syntactic'—which, in this somewhat eccentric use of that word, means that they operate solely in virtue of their physical properties and have no introspectable psychic aspect. Hence, unlike sense-data, our knowledge of them is, at best, topic-neutral.

This sort of externalism is vicious, as is made clear by the following arguments.

1 If the identity—that is, the content—of an experience depends conceptually on what standardly causes it, this being a contingent relation, then one can know what the content of a certain type of experience is only by knowing what its standard cause is. But we could never be in a position to ascertain what the causes of our experiences were for we have, in general, no other access to what objects are around except via experience. We cannot, that is, set up a correlation between experience and cause so as to determine what the experience is seemingly of, for we could not, in principle, adopt the God's-eye viewpoint this would require. This shows that the apparent object of an experience must be something intrinsic to it, not something that stands in a contingent relation to it.

2 Looked at from a slightly different perspective, the same facts render the theory meaningless. The thesis is

X is an experience as of F iff x is that state of the subject standardly caused, through the sense organs, by Fs.

It follows from this that all that we relevantly know about an experience is that it is standardly caused by such and such a stimulus. Such a way of characterising experiences will be contentful only if we have some positive conception of the sorts of thing that can be substituted for F. But as these are features of the external world, they are the sorts of things we can only come to know about by experience. But if experience, considered as a conscious state, has no other properties than that of being standardly caused by a certain stimulus, then there is nothing in the experience which could communicate what the stimulus was like: awareness of such a minimally—or topic-neutrally—known thing could not tell one the nature of its cause. It could do this only if it bore the mark of its cause intrinsically, in some way, and this could be so only if the typical cause was intrinsic to it as object. This is to return to making the intentional object intrinsic to the experience, which is a form of internalism.

5 Externalism, knowledge and perception

Peter Smith has objected to this argument.¹⁵ He denies that the impossible God's-eye perspective would be required before an externalist could know what they were perceiving, and, hence, before we could acquire our concepts. This predicament is avoided by marrying the externalist's causal theory of content with an externalist theory of the knowledge of content. According to an externalist theory of knowledge, a subject knows they are perceiving an F if they are in fact perceiving an F in what is, in fact, a standard way, and they believe they are perceiving an F, the belief being reliably caused via the standard form of perception. Smith says that I might object to the externalist theory of knowledge, but that that is another matter.

An externalist theory of knowledge cannot be used to solve the problem because, uncontroversially, knowledge involves belief and, if my original argument is sound, the subject would be unable to possess the appropriate belief. Knowing that one is perceiving an F involves believing that one is, and believing this involves possession of the concept of F. Supposing F to be a sensible feature of the world, the subject's grasp on F rests, in the end, on experience, by however circuitous a route. (No *simple* empiricism is presupposed here; only that, in the end, our grasp on empirical features of the world rests on the fact that we experience some of them.) Thus, (1) knowledge presupposes belief, and (2) beliefs of the sort we are concerned with

involve concepts ultimately derivable only from experience, therefore (3) that experience is of such a nature as to make possible the possession of such concepts is a presupposition of any conception of knowledge, whether externalist or internalist. But (4) it is the brunt of my argument against the externalist account of experience that it would render impossible the possession of empirical concepts. Therefore, (5) it is not possible to appeal to any account of knowledge to circumvent the kind of objection I make to the externalist theory of experience. An externalist account of knowledge cannot be used to leap over the inadequacies of the externalist treatment of experiential content, given that knowing involves believing, and believing—with the appropriate empirical content—presupposes that experience is able to inform us of the nature of its objects: it presupposes too, therefore, an understanding of experiential content which is consistent with its ability so to inform us. If my original argument is sound, the externalist account of content is not consistent with its possessing this ability.

6 Martin's objections

Michael Martin has made two objections to the argument, one against its 'knowledge' form and one against its 'meaningless' version.¹⁶ Against the former he asks why knowledge of the standard cause of an experience should be the only way of knowing its content. Why could one not instead recognise it by its effects? These effects would, presumably, be the behavioural dispositions that define its content.

This suggestion is answered if one accepts my earlier argument against behavioural theories of perception. But the present argument is meant to be, as far as possible, autonomous, so I shall employ a different argument.

We are concerned with how, on an externalist theory of content, the subject can tell what the content of their experiences is. If this problem is to be solved, that by which one recognises the experience has to be enough to reveal its content. A disposition will do this only if one can—somehow—read the dispositional content in some detail—that is, tell, with some accuracy, what it is a disposition to do. This might be done in either of two ways. One might read it as the 'red-appropriate' disposition, for example. This will be no use in telling one what red is. Alternatively, one might read it in terms of its behavioural detail—that is, in terms of all the things that go to make

up 'red-appropriate' behaviour. Apart from anything else, it is inconceivable what the direct grasping of such a 'multi-track' disposition could be, or be like. But, more seriously, such a thing could not be a general solution to our problem, for reasons that are similar to those that motivated the principal objection to dispositional theories. The content of such a disposition would presuppose a grasp on all those concepts that go to make up an understanding of behaviour in the world. So it presupposes a pretty full battery of empirical concepts. As the problem we are concerned with is seeing how perception could give us access to these features of the world, these concepts would have to be innate. The thesis that we can recognise the contents of our perceptions because we have the innate conceptual ability to recognise the contents of our own behavioural dispositions, seems very implausible.

Martin's objection to the second version of the argument is that, although having a concept of some quality as the standard cause of an experience involves having the experience, it does not involve having a concept of the experience as being as of that quality. Having the concept of the quality depends on having the experience, but this could be merely a causal fact and there not be a conceptual dependence of identifying the cause on being able to identify the experience.

Mere causal dependence is not enough, if the experience is to be the vehicle by which we come to grasp the content. If ordinary naive realism were involved, then one could argue that the experience was 'transparent' and merely made accessible the qualitative nature of the object. But, on current theories, the relation between the quality and the experience is purely causal. If the experience is to be that by which the object is discriminated, and that by which our conception of it is given content, then one must be able to discriminate the experience from others and must have some conception—in however primitive a form—of what it is like to have that experience, so that that can form the core of our notion of the content of the quality: otherwise the experience will not do the job required.

7 The 'straw man' objection

The strongest line of objection to my argument is to claim that it is directed against a straw man.¹⁷ The argument shows that experience *qua* experience—that is, from the viewpoint of the subject—must have its content internal to it. The difficulty comes from the fact that

the externalist need not want to deny this. If this is so then I am wrong (p. 138) to attribute to them in argument (1) that 'one can know what the content of a certain type of experience is only by knowing it as having that cause': and in (2) that 'all that we relevantly know about an experience is that it is standardly caused by such and such a stimulus'. The externalist would not accept that the subject knows about the content of their experience by knowing about the causal relation. Such relations in fact constitute the content but the subject still experiences that content directly. The externalist in fact affirms the following:

- 1 From the viewpoint of the subject, content is internal to experience.
- 2 There is some state of the brain which counts as a representation of an external feature of the world because of the standard causal relation that holds between that brain state and that feature.
- 3 The experience's having the internal content that it does is constituted by the holding of the causal relation between the representation and the external feature. So (1) is constituted by (2). This is a philosophical thesis, so the sense of 'constituted' is such that the constitution is discovered by philosophical reflection or analysis rather than empirical investigation.

The mind-brain identity theorist will add

- 4 The experience is identical with the brain state.

The problem that one feels this theory faces is to explain how the existence of an external causal connection can be experienced as an internal content. This problem is particularly clear if one accepts the identity theory, for then that which is external and that which is internal are external and internal to the same thing. Without including the assertion of identity between experience and brain state, the experience might include more than the brain state and, therefore, something external to that brain state could, without paradox, be internal to the experience. This would still not actually remove the problem, for if the experience were hallucinatory it could occur in the absence of anything external and so there would still be a *prima facie* problem about how its content could be constituted by something in the external world.

The arguments of the preceding pages are directed against a particular answer to the question of how an external cause could constitute an internal content subjectively. That answer is that they could do so only by the subject's recognising it as the inner state that does indeed have that standard cause; and the arguments above do show that that is not an acceptable answer, for one could never be in a position so to recognise it without already being able to recognise the internal content first. But what other possible answers are there? There was a rationale behind my thinking that the strategy already refuted was the only one available. This was that if a kind of state of affairs *A* is wholly constituted by some kind of state of affairs *B*, then anything going on in *A*-terms ought to be explicable in *B*-terms. When this is a philosophical thesis, the explanation should not merely be causal, but something more nearly analytic or conceptual. So we are owed an explanation of how a distal stimulus could be experienced as subjective content. That one recognises it by its effects on one seems about the only plausible candidate, and this is the strategy shown to be impossible. So if the experiencing of external cause as internal content does not occur in virtue of the induction I have shown to be impossible, how does it occur?

In fact the principal physicalist accounts deny that the subject recognises experiential content at all, in any real sense. A behaviourist or a functionalist says that the typical cause constitutes the internal content only by determining the dispositional properties that constitute the experience. One might still want to raise the question of how the subject recognises this content. In reply to Martin I have already argued that they do not do so by identifying the disposition. The standard behaviouristic answer, however, does not have a proper role for recognising the content. It says that being aware of the content of one's own experiences is essentially a matter of being trained to insert verbal responses, and, perhaps, other relatively sophisticated responses, into one's natural behaviour. One does not respond as a result of recognising some content one has; there is no more to having and recognising content than the response. It is always worth reminding oneself how utterly implausible this kind of theory is. Rather than accepting that there is a genuine phenomenology, of which one can be introspectively aware and which one can, therefore, report, the phenomenology and the introspection are no more than the disposition to make the report. This account simply abolishes the first-person perspective (which is a much stronger position than denying it priority

or autonomy). This is most clearly brought out in the case of hallucination. In an hallucination the subject seems to have a grasp on some sort of private content, however interpreted. On the dispositional account, all they have is a disposition to behave and report. This must not be mistaken for the theory that the hallucinator's subjectivity consists in them being *aware* of having certain dispositions: it consists simply in their *having* the dispositions. This is why I say it abolishes the first-person perspective. If one fills out the account by bringing in awareness of the dispositions, then the externalist is back with Martin's theory.

This behavioural theory is not merely counter-intuitive; it renders the origin of our recognitional abilities a mystery. It takes as given the community's understanding of language, and pictures the community as training linguistic response into the learner's behavioural repertoire, without having any account of the learner's experience—how it is with them—which would make their uptake of the teaching intelligible. If each individual subject lacks a pre-linguistic intuitive grasp on content, why should they come to associate words with any particular experience—why should the learning game be played at all? It will be said that this question reflects an individualistic conception of experience: the behavioural usefulness of language gives an evolutionary explanation of why it should develop. This response is not satisfactory. The usefulness of something does not explain *how* it can develop, and, however much the terms in which the subject interprets experience may depend on social influence, sensory experience is something the *individual* has, and language must latch onto this thing that belongs to the subject *qua* subject.

Armstrong tries to make the dispositional theory more plausible by bringing in awareness of the dispositions.¹⁸ He identifies dispositions with causal pressures towards behaviour which are felt and identified by other parts of the brain: this is what consciousness and introspection are. But this is no good as an account of perceptual content. Armstrong fails to distinguish between the immediate causal output of a brain process—which will be a small electric charge—and its causal significance for the behaviour of the organism as a whole, which depends on the rest of the body, inside and outside of the brain. Simply 'feeling' the pressures within the brain will tell one nothing of their dispositional content unless one is aware of the context: but one has no way of knowing that context for it is precisely 'feeling' these pressures which is supposed to constitute consciousness of one's perception of

things external. Armstrong's account of consciousness is vicious in just the way I have been attacking.

In general, it looks as if any identity theory will have difficulty with externalism. The brain state is purely internal to the head and its causal properties are limited in range, as I have just argued. If this is identical with experience-with-content, then content must be somehow internal to the brain state, which is just what seems impossible. A non-identity version of functionalism does not have the problem in this form, for there is no question of the *disposition* falling short of the required content, in the way that do the immediate causal powers of the brain state. There are still formidable problems. If one wishes to make a gesture towards acknowledging the first-personal perspective, by saying that it consists in being aware of one's dispositions, then there is the problem we raised for Martin. In fact, such an awareness could only be further dispositions, and, as I have argued, this abolishes the subjective view altogether.

8 Macdonald's weak externalism

Cynthia Macdonald has presented an argument which, if sound, saves the identity theory from my objection.¹⁹ Her claim essentially is that, considering kinds of mental state, their content is defined externally, but as tokens, the content is internal. So, she argues, externalism and the mind-brain identity theory are compatible. That they are not compatible is argued on the simple grounds that if something *external* is constitutive of the content of a mental state, then, *ex hypothesi*, the mental state is not identical to a purely *internal* brain state. Macdonald's reply is that something external may be essential to the characterisation of an entity without being constitutive of it. For example, an offspring is essential to make someone a father, but the offspring is not in any way constitutive of the man who is the father. Similarly, the causally related external objects are necessary to make the brain states into mental states, but it is still the brain states that are the mental states. This she calls 'weak externalism', because the external object is not constitutive of the particular mental state.

As the analogy with fatherhood shows, Macdonald's point is a very simple one—so simple as to make one wonder how a collection of astute philosophers, such as Burge and McGinn, should have omitted to consider it. The point is, of course, that they assumed that content is

constitutive of a mental state, and not merely something that qualifies something as a mental state: it is more like the relation of child to family than of child to father. But why should this be so? What sort of fact is it that determines content to be constitutive rather than a qualifying characteristic?

In order to see the answer to these questions, one must consider the relation between externalism and consciousness. There can be an ambiguity about the term 'externalism'. At the start of this discussion I gave as the criterion for externalism that the content be essentially object-invoking. There is, however, another criterion not far in the background, namely that the content must transcend the reflective or 'introspective' awareness of the subject. For someone who thinks—like a traditional empiricist—that the mind can be reflectively aware only of 'representations' and not of external things, these criteria coincide; for someone who thinks that the mind sometimes really reaches out to objects, they do not. Reflection transcendence is the more important criterion for the philosophy of mind in so far as it is concerned with conscious states.

Perceiving a father does not, *ipso facto* make you perceive him *as* a father, because it does not include perception of a child. Perceiving a family does involve perceiving a child. Consciousness of a mental state *as* a mental state—that is, the sort of consciousness that someone has of their own mental states—necessarily includes consciousness of the content: without that there would be nothing of which to be aware. So from the perspective of reflective consciousness, content is constitutive of the mental state: it is like perceiving a family rather than like perceiving a father, in that all the relata must be included in the grasp and not merely exist outside the scope of the awareness. This does present obvious problems for the mind-brain identity theory. If a given mental state and a given brain state are identical, and if it is possible—as it is—to be aware of one's mental states as mental, then there must be some way of being aware of the brain state as a mental state. (This would not, of course, be by viewing the brain from outside, but by the kind of access the subject themselves may be supposed to have.) We have already agreed that awareness of a state as mental involves awareness of its content, but no awareness of a brain state *per se* will involve awareness of any external object standing in a causal relation to it. Considered as a state of which we can be reflectively conscious, therefore, a mental state cannot be identical to a brain state.

The only possible way round this argument would be if we could be aware of the brain state as *the state standardly brought about in us by a certain cause*, or some similar formula. But this brings us back to the original argument against externalism. Macdonald's 'weak externalism' does not circumvent the problem. The second premise of the original argument is justified—externalism about perceptual content, in the form which physicalism requires, is demonstrably false. So reductive physicalism about perception is refuted.

9 External to the mind or the head?

The above discussion shows that experiential content cannot be external to experience. There is, however, an unclarity about what externalism, as usually discussed, is making content external to. Externalism holds that content is conceptually determined by something external to the body. Does it follow from this that it is determined by something external to mental states as they are 'folk psychologically' conceived? The traditional mind-brain identity theorist is identifying the mind of folk psychology with the brain, so the answer in that case is that if the content is external to the body it is, *pace* Macdonald, external to the mental state as well. Modern cognitive scientists may not have the same view. Proponents of the *syntactic theory of mind* claim that the 'narrow content' that is internal and which causally drives our behaviour is innocent of any semantic properties. 'Broad content', which corresponds more to our folk psychological conception of what is in our minds (and if it did not we would presumably not have such a notion; it exists to compensate for the fact that the narrow content in the head seems to leave out something pre-theoretically important) includes semantics and the relevant bits of the external world. This raises the question of whether we can be introspectively aware of broad content. This presses us to a further development of the question of the relation between folk psychology and externalism. The term 'folk psychology' is generally taken as referring to our normal practice of explaining behaviour in terms of beliefs and desires, and, hence, to beliefs, desires and other propositional attitudes. Because of the connection with *explanation*, these propositional attitudes have tended to be treated as theoretical entities. But there is another, less abstract, version of the lay conception of mind. According to that, our ordinary mental states are things of which we can be reflectively or

introspectively aware. One conception of externalism would be, therefore, that the contents of mental states are *reflection transcendent*—that is, they are not the sorts of things that we can capture by becoming aware of, or self-aware concerning, our mental states. The habit of calling externalism ‘anti-Cartesian’ suggests this kind of externality. So do some of the standard examples. No amount of reflection on one’s own experience will tell one whether the ‘water’ in one’s environment is H O or XYZ, or whatever else, microscopically. If applied to basic perceptual content, such a theory is obviously absurd; I can obviously be reflectively aware of how it seems to me. In so far as the argument of these sections concerns the theory that content is external to experience, considered as something one can be reflectively aware of, then argument is hardly needed to dismiss such a theory. When it comes to experience in a folk psychological or lay sense, all that the externalist could mean is that there is a conceptual connection between some inner ‘representation’—in the sense in which the syntactic theory of mind employs that term—having a standard external cause and the possession of a certain content internally by an experience, in the ordinary sense. There are two things to say about such a theory.

The first is that experience in the lay sense, and meaning basic phenomenal content, is not naturally treated as part of broad content. It is so natural to think of *qualia* or sensations as bed-rock and internal that even Churchland tries to treat *qualia* as brain states. (His failure was recorded earlier in this chapter.) The second objection is that the theory does not work, if the same inner representation that operates in a perception could produce a corresponding veridical hallucination. (That it could is defended in the argument which I begin on p. 151.) If there could be such an hallucination, then it would follow that the inner state was sufficient to constitute (rather than cause, from a physicalist perspective) the appropriate experience, given the rest of the brain context. Veridicalseeming hallucinations have the same introspectable content as perceptions, whereas if the experience really supervened on the inner representation and the outer object together, the inner representation alone would not be enough. It might be tempting to argue at this point that, although the inner representation does have a certain ‘feel’ to it, this is somehow inchoate until it is interpreted as relating to the object, and this can occur only if, in fact, it is standardly caused by the object. This argument, however, falls to the original argument, for,

if the experience is inchoate, it does not reveal the object and we have no way of identifying it.

There is a general moral behind this difficulty with the externalist treatment of experience. This moral is that any case of what I shall call *essential externalism* is invariably vicious, whatever mental content it is applied to. Essential externalism can be explained by contrasting it with accidental externalism.

When Kripke says that someone—call him Jones—could refer to Richard Feynman, even though he knew nothing about him, on the strength of hearing others use the name, he could be taken as suggesting that Jones could think a thought the content of which was—in part, at least—external to his mind, in the sense of being beyond his reflective grasp.²⁰ But those who know Feynman could think the same thought internally, in the sense of being completely aware of its content; the externality of Jones' thought is accidental to its content. Indeed, it is parasitic on cases in which it is not external—someone must, sometime, have known Feynman properly for the name to have been coined. By contrast, an essentially externalist account is one that denies that a certain kind of content ever could, in principle, be open to reflective scrutiny. Any account of a kind of content that says that its relation to the relevant mental state is purely causal—as was the case with the analysis of experience I was attacking—is essentially externalist. So that, just as Jones when thinking about Feynman contingently cannot rehearse reflectively to himself the content of his thought, if externalism is essentially correct for certain classes of contents, no one could ever rehearse reflectively to themselves what these contents are. If this is the case, none of us ever reflectively understands what is happening in our minds when we employ these contents. As an account of perfectly ordinary thoughts and other cognitive states, this is absurd.

(Whether current theories of the semantics of natural kind terms are essentially externalist is unclear to me. Discovering that water is H O may be parallel to discovering who Feynman is, thus rendering the semantics of 'water' only accidentally external. What is clear is that this sort of extension of knowledge does not essentially affect experience. There are those who emphasise the plasticity of experience in the face of theory and conceptual change; but, whatever truth there may be there, it cannot imply that, if you get the right theory, H O and XYZ will look different. Even for ordinary thoughts, it is not clear whether *de re* semantics are viciously externalist. If the

Perception

semantics of ‘hydrogen’ and ‘oxygen’, as well as those of ‘water’, are externalist, and so on, maybe one is in a predicament not very different from being in Searle’s Chinese room. But this is a matter too large to pursue here.)

CHAPTER VI

The Revised—and Successful— Causal Argument for Sense-data

1 The argument stated

The revised causal argument is a combination of the original causal argument and the argument from hallucination.¹ It can be stated in the following way.

It is clearly true that

- 1 It is theoretically possible by activating some brain process which is involved in a particular type of perception to cause an hallucination which exactly resembles that perception in its subjective character.
- 2 It is necessary to give the same account of both hallucinating and perceptual experience when they have the same neural cause. Thus, it is not, for example, plausible to say that the hallucinatory experience involves a mental image or sense-datum, but that the perception does not, if the two have the same proximate—that is, neural—cause.

These two propositions together entail that perceptual processes in the brain produce some object of awareness which cannot be identified with any feature of the external world—that is, they produce a sense-datum.

Proposition (1) claims that there is some state of a subject's brain which is sufficient for their having a particular type of experience. For example, suppose that they face a brown table stood against a green wall, with their eyes open and their sensory mechanism functioning normally. There is light around and a certain brain state is induced in them in accordance with certain partially known mechanisms, and they have the experience of seeing the table against the wall. Now suppose that their brain is fixed, in the relevant areas, in this state, and the table

and wall hidden from them. Such a fixing of brain state is conceivable, not only in the sense that it is logically possible; it is empirically possible, in the sense that we may conceive of ways of doing it which involved no breach of natural law. This is compatible with its not being practicable and even with our having reasons for thinking that it might never be practical. It might even happen by a remarkably improbable coincidence (given quantum theory) or there might be a drug which held constant the relevant areas of the brain. Proposition (1) asserts that if such fixing took place the subject would continue to have an experience exactly similar to that of seeing the table and wall—they would ‘see’ a brown table-shape against a green background—even though the table and wall were covered and removed, so that they no longer acted on them.

Because there is fairly general agreement that our mental life is strongly tied to what happens in the brain, there is little inclination to deny proposition (1). Even those who are externalists about perceptual content tend, I think, to accept the sufficiency of the brain for the production of experience, phenomenally conceived. They might argue that what that experience is like depends on what is *standardly in fact* the brain state’s external cause, but they would still accept that the brain state would sustain an experience of that kind in the deviant instance when the standard cause is missing. But even if this is not so, it does not matter, for we have already seen that externalism about phenomenal content is false.

There seems, furthermore, to be a reasonable amount of empirical evidence for (1), and it also explains why veridical-seeming hallucinations should occur.² If it were not for the fact that perceptual processes, however stimulated, were sufficient to generate experience, it would be a mystery why such hallucinations should occur.

Nevertheless, (1) can be attacked, if rather desperately. Because this is a desperate strategy, and because, therefore, attacks on (2) are more common, I shall return to (1) after considering the rather popular objections to (2).

2 The attack on (2): the disjunctive analysis

In order to reject (2) and defend naive realism it is necessary to deny that the relevant brain process produces more than a bare act of awareness in the case of normal perception, whilst allowing that the same process produces an internal object or content when artificial

stimulation produces an hallucination. 'I agree', the direct realist might say, 'that being in a certain brain state is a sufficient internal causal condition for experiencing certain sense-contents, in this case a table and a wall, but I believe that we must give a separate account of hallucinatory or artificial experiences from those that we give of normal experiences. When the objects are removed and the brain state frozen, this is hallucination, and the subject 'sees' or 'has' images of the table and wall, but when the objects are the cause of the brain state, the subject sees them themselves, directly.'

Don Locke adopts an argument rather like this in his attack on those who argue for sense-data on the ground that perception and hallucination are qualitatively indistinguishable. He calls their argument 'the argument from qualitative similarity', and says that this presents us with a choice:

We can either allow that things that are, or at any rate could be, qualitatively alike are nevertheless ontologically distinct, or agree that we never perceive external objects. There seems no good reason for accepting the second alternative.³

In so far as Locke is merely pointing out that qualitative similarity does not entail ontological similarity, he is only using the same argument as was used in Chapter III against the argument from hallucination. The principle that qualitative similarity does not entail ontological similarity is accepted by most sense-datum theorists, for representative realists assert that data and physical objects share some properties. (Berkeley, of course, dissents, on the grounds that 'an idea can be like nothing but an idea'.)⁴ It might also be true that we ought to give a different account of hallucinations as they in fact occur (for example, to schizophrenics, drunkards and drug-takers) than we give of perception, but those who suggest that such hallucinations need separate treatment presumably do so because they believe that normal hallucinations are very different from perceptions in their immediate causes—that is, that they involve the activation of different brain mechanisms. But in the present argument just the same brain state as is the immediate cause of the perception is the immediate antecedent of the corresponding hallucination. If the mechanism or brain state is a sufficient causal condition for the production of an image, or otherwise characterised subjective sense-content, when the table and wall are not there, why is it not so sufficient when they are present? Does the brain state mysteriously know how it is being produced; does it, by some extra

sense, discern whether the table is really there or not and act accordingly, or does the table, when present, inhibit the production of an image by some sort of action at a distance? That such *ad hoc* hypotheses are implausible is the rationale of (2). Surely, difference of ontological status cannot plausibly be attributed to the sense-contents in question if they are the upshot of the same proximate causes. This conviction, however, is not universal. Dretske in *Seeing and Knowing* actually considers a case like the one I hypothesise in which the hallucination is caused by precisely the same brain process as causes the normal perception, yet treats it as just another form of the argument from hallucination: he therefore deals with it as Locke deals with hallucination:

The pattern of this argument, the Argument from Hallucination, is obviously fallacious. I may not be able to distinguish between S's handing me a genuine one-dollar bill...and S's handing me a counterfeit one-dollar bill... If the counterfeit is good then these events may be indistinguishable. Surely, however, we cannot conclude that because I am being handed a counterfeit bill in the one case I must be receiving a counterfeit bill in the other case?⁵

Dretske shows no recognition that it might make any difference whether the immediate causes are the same.

I am aware of only two philosophers who face the case where hallucination and perception operate through precisely the same brain process and see the need to defend the view that nevertheless the products of that process are different in the two cases. The philosophers are Hinton and Pitcher.⁶

The general principle lying behind (2) in the original argument could be expressed in the slogan 'same proximate cause, same immediate effect'. Call this slogan 'S'. Pitcher accepts S, but in a form that would not sustain (2). He says:

This principle S although it may be true for every cause-and-effect pair, is not true for them *under every description*.⁷

The idea is that though the immediate effect of a certain brain state may be always 'seeming to see something red', and thus, in that respect, always be of the same type, they may also be of different types, in that, sometimes, the immediate effect will fall under the description

'hallucinating something red' and, sometimes, under the description 'seeing something red'.

The general principle that immediate effects of the same type of proximate cause need not be of identical type under all descriptions is quite sound: indeed, no two *different* events could be of identical types under all descriptions true of them, unless they were corresponding events in mirror universes. More concretely, if someone strikes identical nails into identical walls with identical hammers with identical force, the effect in one case might be described as 'the picture's being hung' and in the other as 'the gas pipe's being severed', each description applying to only one of the effects. The description common to both effects will concern a nail moving a certain distance into a wall. It is plainly not arbitrary, however, what the common description is. If it were sufficient for the satisfaction of S simply that there be *some* common description then it would become vacuous. Suppose that, despite the qualitative identity of hammers, nails, walls and force of blows, in one case the nail penetrated one inch and, in the other, two inches. Intuitively, this would involve an infringement of S, for the same cause would have had relevantly different effects. Nevertheless, the effects are similar under some descriptions; for example, the description 'a nail penetrating a wall'. Even if in one case the nail had burst into flames the effects would have been identical under the description 'something's happening to a nail'. It is often said that any two things are similar in some respect: if this is true then any effects of some one type of cause will be identical under some description or other.

To save S from vacuity it is, therefore, essential to say something about the descriptions under which the events must be similar: some restrictions must be imposed. Intuitively, it is not difficult to do this. The different descriptions which applied to the nail applied in virtue of features of the situation which were more remote from the immediate cause—the hammer blow—than was its immediate effect, the nail's movement. They refer either to context—the presence of the picture—or to a further effect—the piercing of the pipe. Intuitively, S applies to the most specific and immediate characterisation of the effect. Thus the communality of a generic description such as 'entering the wall' will not satisfy S if one nail entered by one inch and the other by two. One requires the kind of exactness that might answer to, or follow from, a natural law, and be susceptible to quantification. The difference between the nails penetrating one inch and two parallels what Pitcher

says about the mental state caused by the brain process. He says that the same brain state will always cause a 'seeming to see something red' but that sometimes it will do this by causing a genuine seeing and sometimes by causing an hallucination, where these two states are essentially different. We do not have a genuine case of S where the effects brought under the same description are analysed as having radically different structures or component elements. A sense-datum theorist has no problem with this, for they say that the same brain state causes one kind of effect—the having of a particular kind of sense-datum—by doing which it causes either a seeing or an hallucinating, depending on the further causes and circumstances. But for Pitcher seeing and hallucinating are not to be analysed into a truly common element plus differing extra features.

It seems, therefore, that when one realises the need to apply S to certain sorts of descriptions, Pitcher's argument will not work. Hinton does what is required and simply denies S. He says that it rests on 'dubious and arbitrary metaphysical beliefs...about effects of causes'.⁸ He presents an explicitly disjunctive analysis of such generic concepts as 'experience' and 'seeming to see'. Such expressions are taken not as referring to something common to both perception and hallucination but generically to a disjunction of them both. Thus he deals with the final outcome of the process as follows:

The impulse reaches certain specified structures, and then what? My continuation was 'and then one perceives a flash of light or has the illusion of doing so, as the case may be, according to the nature of the initial stimulus'.⁹

He is quite aware of the motives that give rise to belief in S.

But it is natural to make some such retort as this, that what happens cannot depend on the initial stimulus; what happens next must be the same, whether the initial stimulus was light striking the retina, or an electric current passing through the retina, or whatever it was.¹⁰

But he rejects this 'natural retort' on the grounds cited above, namely that it rests on arbitrary and dubious metaphysical beliefs. Hinton holds that S is a sound principle when applied to causal laws that relate physical or public events, but that it is mere prejudice to extend it to psycho-physical connections (or, for that matter, to mental-mental ones). His argument is that we know that S applies to physical laws

because we can in those cases identify the effect independently of identifying the causes—this is part of the publicity of physical events—and have thereby been able to establish empirically the truth of S in these contexts. But we cannot do this for mental effects; for example, it follows *ex hypothesi* that one cannot tell the type of hallucination that we are considering from a perception without knowing the causal ancestry of the experience: the experience itself does not reveal which it is. Therefore we are not compelled to apply S in the mental case, for we have not identified the nature of the effect independently of the causal context and verified that S applies.¹¹ However, this is an adequate argument only if there are no general considerations in favour of S—only if, that is, S relies on empirical proof in each type of context.

First, we can concede that one is under no compulsion to accept S, if by compulsion one means logical necessity. Second, we can concede that S is more incontrovertibly established in purely physical contexts. It seems, however, that there are good reasons for extending it to psycho-physical contexts and, indeed, for adopting it as a perfectly general principle. One of the reasons for applying it in psycho-physical contexts has already been alluded to: how would the brain state know when it is required to produce an image to act as understudy for a genuine perception, and why should it bother to do so, as the hallucination serves no purpose except to deceive? If direct realism were correct for perception, then one would not expect (1) to be true; that is, the same process should not be suited to producing an hallucination. The second reason is a variant on this. Unless S applies in the psycho-physical case, the existence of the hallucinations becomes a mystery. Given that Hinton is conceding that hallucinations could be produced by stimulating just those brain states involved in perception, how are we to make sense of why this should be so if it is not by thinking of hallucinations as cases in which the state is activated and performs its normal function—that is, has its normal causal upshot—in an abnormal context? Otherwise the production of hallucinations in this way would seem to cast the brain in a role something like that of a Cartesian demon, producing an effect specialised solely to the context of deception. In order to make sense of why hallucinations should be generable from perceptual processes, I suggest that we have to be able to make sense of ‘what the brain state does in both cases’—that is, of a common element not analysable into Hinton’s disjunction. Thus we face a choice between

accepting the radical unintelligibility of why there should be hallucinations in these contexts or of accepting that S is applicable here.

It seems, therefore, that S is sound (contra Hinton) and sound in a strict enough sense (contra Pitcher) to justify (2).

3 Anti-realist intuitions

It might be argued that this response to Hinton is excessively *realist* about experience. I talk as if it is some sort of entirely objective matter of fact that the mental state produced by the brain state is thus and so, whereas, it might be argued, psychological states are a matter of interpretation and construction. Davidson's and Dennett's theories of belief-desire psychology are instances of the rejection of a hard-core realism about mental states and its replacement by a theory according to which the mental is constructed by our interpretative practices.¹² It would be wholly bizarre to suggest that the very *existence* of sensation depended on such practices, but perhaps less so to claim that its exact structure and logical features depend on what we make of it. It would not, therefore, be some independent question of fact whether the psychological state that was our perceiving something were identical in kind with the corresponding hallucination, but a matter of how our language or theorising about perception and hallucination structured them.

We are now onto a theme that we have met before, when discussing colour in Chapter III, and will meet again when discussing intentionality, namely the dependence of mind on language or 'grammar'. There seem to me to be at least two objections to its present application. First, the hybrid mixture of realism and interpretationism about sensations is not plausible in this case. It is not that the idea that brute sensations which exist independently of thought might be given structure by thought is, on reflection, implausible. Rather it is that, in this case, the direct interpretations of perception and veridical hallucination do not seem to be different. In both cases, we structure the experience as being as of an external reality; it is only at a second level that they are distinguished, when the hallucinatory experience is discounted. This defeasibility seems too remote a feature to make any difference to how we characterise the direct product of the brain process.¹³

The second objection is like one we found when discussing the ‘grammatico-dispositional’ approach to colour. There we found that it does not show a way of coping with the facts of ‘illusion’. When a white object looks red, for example, the direct realist has a problem in explaining how red fits into the story. This is just as much a problem if one’s realism comes from the ‘grammar’ of language as it is if it is taken to be simply an expression of common belief. This brings out a general limitation with the disjunctive theory. It is natural to think of the disjunctive theory as constituting an attempted account of the difference between veridical and non-veridical perception. In fact, it can only apply to the contrast between perception and hallucination. In other words, it leaves untouched all those phenomena categorised as *illusions*, which are non-veridical *perceptions*. If one were to apply the disjunctive approach to illusions one would play straight into the hands of the argument from illusion. The disjunctive analysis allows that hallucinations involve awareness of something subjective, but that perceptions do not. If all nonveridical perceptions were treated in the same way as hallucinations, then every case of something not looking exactly as it is would be a case in which one was aware of some kind of subjective content. Only perfectly veridical perceptions would be free of such subjective contents. This is exactly the situation to which the argument from illusion leads, and we have already seen how this naturally collapses into the view that there always is subjective content and never direct awareness of the external object.¹⁴

The conclusion is that the disjunctive analysis is radically implausible, both in realist and not-so-realist versions; and that even if it were not, and could cope with the causal-hallucinatory argument, it is seriously incomplete, focusing our attention back onto the argument from illusion.

4 Defence of premise (1)

Given the failure of the attacks on (2), we must return and see whether (1) is more vulnerable than it initially seemed.

One approach would be to challenge the straightforward realism about sensation that seems to be implicit in (1): it seems to be taken as a brute fact that brain states cause sensations or experiences, on a par with the fact that they cause further brain states. The discussion at the end of the previous section effectively covers this option. An anti-realism that makes sense-contents entirely dependent upon

interpretation, denying any 'brute' element, is wholly implausible; and a hybrid version has already been discussed.

Proposition (1) asserts that the relevant brain process always gives rise to a qualitatively similar experience, whether or not there is an ontologically common element. Hinton, as well as denying that there is an ontologically common element, also denies that there is any strict qualitative similarity. Hinton does not deny that perceiving and hallucinating may be qualitatively indistinguishable. He does deny that it follows from their indistinguishability that they are exactly similar. He thinks that by distinguishing inference on the basis of perception from the given nature of perception itself, he can avoid saying that the event of perceiving and that of hallucinating have an exact qualitative similarity, for it is only that we are inclined to judge that they have.¹⁵

The first thing to say about this claim is that it has a point only as a prolegomenon to the disjunctive analysis: only if the brain process is followed by different effects in the two cases will it be relevant to deny qualitative similarity between the effects. The second is that it is *per se* implausible. One might judge qualitatively different experiences to be similar if one were hurried or slapdash, but if they appeared similar however closely they were introspected, then the suggestion that they were really different seems implausible and arbitrary, if not vacuous. Is Hinton suggesting that a close enough inspection would always show up a difference? This could be so only if S were denied, and S has been adequately defended already.

Given a sensible realism about experience, there is only one way that I can see in which proposition (1) can be denied. The defence of (2) shows that one is obliged to accept that the relevant brain state produces the same effect in both hallucination and perception. If this is not to involve something like a sense-datum in the case of perception, neither must it in hallucination. For the naive realist, the proper contents in perception are provided by the sensible features of external objects, and in hallucination these are missing, so the hallucination cannot seem like the corresponding perception. Under these constraints, what could be said about the factor common to both hallucination and perception?

The best that can be done is to say that hallucination consists of a kind of imaging which is phenomenally distinct from perceiving, and that this is the factor common to both. This imaging could be thought of as being like the imaging one does when one tries to conjure up what

something looks or sounds like. As its connection with such thinking suggests, imaging of this kind is a kind of thinking—a sort of image-conceptualisation. This explains its role in perceiving. Perceiving will have two facets. First, there is the direct consciousness of external things, second there is imaging by which this is appropriated or grasped, and this second depends entirely on the brain: in so far as a feature of a perception is not imaged one has no reflective awareness of it. Thus one is only aware of perceiving some feature of the world if, in addition to consciousness falling upon it, it is also imaged. This would explain misperception of certain kinds. If I remove my glasses, my imaging becomes blurred and I am only able to assimilate confusedly what is in my bare perceptual consciousness.

I think this is the best theory that the naive realist can put forward whilst allowing a univocal role to the perceptual mechanism and brain. There are, however, several serious difficulties.

1 The imagist approach to the common element denies that direct stimulation of the parts of the brain involved in perception will give perceptual-seeming experiences. It would be enough for the causal argument if stimulation of the rods and cones in the eye, and, hence, of the process up to and including the brain, caused a veridical-seeming experience. That this should not be the case seems both very implausible and, in so far as there is evidence, empirically false.

2 On this theory, when a white wall looks red one is only imaging red, in a sense of ‘imaging’ in which an image is introspectively distinguishable from perception. But, whatever the situation may be for hallucinations, ordinary misperceptions can often seem to be totally veridical.

3 If the physical process is responsible only for the imaging, there appears to be no principled account of the relation of bare perceptual awareness to any physical process. This seems odd in itself, but the situation is worse when one takes on board the fact that it is the imaging that determines the exactness of the final experience. The bare perceptual consciousness could possess an acuity far beyond what we are ever fully conscious of; we simply do not notice because we cannot image appropriately. Deterioration of vision could be deterioration of our imaging capacity—no bare perceptual change need be invoked. The theory is now reminiscent of the claim that the brain does not cause experience, but only edits most of it out—without a brain the mind would be overwhelmed, and would see the universe in every grain of sand—or, at least, in all its microscopic

detail. There may be an important truth about the role of embodiment in this theory, but it is not helpful here, because the direct perception of the fullness of things would not correspond to the bare perceptual element in naive realism.

5 Conclusion

If this chapter is correct, then there is a subjective element in all perception, for which a brain state is a sufficient cause, and which contains all those phenomenal features that we are familiar with in perception. What is now to be discussed is whether this conclusion can be accepted without also accepting a version of the sense-datum theory. The conclusion does, indeed, appear to be nothing other than the sense-datum theory, but there are ways of construing these subjective contents which are meant to weaken this conclusion, and these must now be considered.

CHAPTER VII

The Intentional and Adverbial Theories

1 Motivation

The conclusion at the end of the previous chapter was that perception does involve a subjective phenomenal content of a kind normally identified with the sense-datum theory. The purpose of this chapter is to see whether one can accept the existence of this content whilst giving it a different interpretation from a sense-datum one. So the argument henceforward takes for granted the falsehood of the disjunctive analysis and accepts the existence of a phenomenal element common to perception and appropriately caused veridical hallucination, and is concerned to see whether there is some way of construing this common element which can save some kind of naive or direct realism—or, at least, to refashion a theory in some way closer to naive realism than is the sense-datum theory.

Resistance to the sense-datum theory is inspired mainly by the fear that such data constitute a *veil of perception* which stands between the observer and the external world. This threatens to engender scepticism, or even solipsism, and runs counter to good, plain common sense. If the veil cannot quite be denied—as the conclusion of the last chapter seems to suggest—perhaps it can be made less heavy and opaque. Perhaps it is not so much a screen between us and the world as the means by which we perceive external reality: or perhaps its ontological status can be made so flimsy that it is too diaphanous to resist our penetrative gaze.

Attempts to draw back from the abyss of the sense-datum theory at the last moment tend to centre upon the ideas that sense-contents are *intentional objects* and, hence, *intentionally inexistent*, or that they are not *objects* at all, but merely *adverbial modes of sensing*. Sometimes the notions of intentionality and adverbiality are treated as if they were

complementary aspects of the same theory and sometimes as if they were autonomous—or, at least, as if it were one or the other of them that was doing all the work.

2 Intentional objects and psychological states

The crucial thing about intentionality is that when an act takes an intentional object that object is *intentionally nonexistent*. I can think about unicorns when no unicorns exist. 'Inexistence' flags the idea that something is both a genuine object of an act and yet unreal. It has to be distinguished from ordinary *non-existence*, for an occurrent (and existent) mental act is not constituted simply by the non-existence of something; how does the non-existence of *that* thing rather than of anything else make the act the one that it is? So 'intentional inexistence' is more of a name for a problem than it is a solution to one.¹ Nevertheless, we do believe that thought about a unicorn is *of* a unicorn without there being any unicorn; and this does not mean that there must be a replica of a unicorn in the mind which is what we are thinking of in its place: there is, no doubt, a representation of a unicorn in the mind when one thinks of a unicorn, and that representation may, sometimes, be a kind of replica, but that is not what we are thinking *of*, rather it is what we think *with* or *by*.

If this intentionalist pattern of thought were applied to perception, the fact that one could seem to see or hallucinate a red patch without there being a red thing in the external world would not force one to say that one were instead aware of a red thing in the mind: hence sense-data would be avoided.

There are two difficulties with this as a solution to the problem of the nature of the common element in perception. The first is more of a preliminary reservation. It is true that thought of a unicorn does not involve the existence of a unicorn nor of a sort of mental replica of one. Nevertheless, there is a difference between the logical features of intentionality and the psychological reality of mental acts with that feature. The thought itself will, in the standard case and if it is introspectable, have a psychological vehicle. This will probably be verbal, and will involve rehearsing a sentence to oneself; ('saying in one's heart'). There is nothing nonexistent about this vehicle; although mental images are, in various ways, problematic, they are not so in the kind of ways that their *objects* are problematic. The image—verbal or otherwise—is not nonexistent, even though what it is of or about is

inexistent. Appeal to intentionality does not, therefore, solve—or, perhaps, even help to solve—any problems there may be about the ontology of the images themselves. One could say that Henry VIII is inexistent with respect to (in a non-literal sense, exists in) Holbein's portrait of him, but the pigment of the painting exists in a wholly unproblematic way. The relationship of mental object to psychological vehicle might be conceived of as analogous to this.² This should lead us to be cautious about attempts to appeal to intentionality to explain the *psychological* reality or the phenomenology of mental states, as opposed to their logical features. In the case of thought, this may not be an absolute distinction, though it is definitely a real one (as the role of 'saying in one's heart' shows): in the case of perception the concretely phenomenological looms larger and is more central. This leads naturally to the second point.

The appeal to intentionality tends to play down or ignore the difference between perceptual experience and other kinds of mental activity where the relevance of intentionality is less controversial. The object enters into perceptual experience more thoroughly than into other kinds of mental act. If you are radically ignorant of the physical properties of an object it is no use trying to find them out by thinking of it, hating it, loving it or desiring it; you must try to come by a perceptual experience of—or as of—it. For example, it is the blind person whose conception of red is most defective, not someone who cannot hate, love or desire it: the tone deaf, who cannot care about music, miss something about the nature of sound, but not anything as fundamental or complete as those who are deaf *simpliciter*. It can hardly be disputed that experience reveals the nature of objects—at least as they are conceived to be in our naive conception of the world—in a way that other mental states or attitudes do not. They are characterised by the fact that they do not require the presence of the object in question; whether an object is there when you think of or desire it is immaterial to the phenomenology of thought and desire as such. They are essentially acts tailored to the absence of their objects. This is shown in the fact that the experiential differences between other intentional states—between fearing, loving or desiring, for example—do not consist in any difference in the manner of the presence of the object, but in the manner of the subject's response to the object. The object is present purely intellectually, that is, as an object of thought, in all these cases. But in sensory experience the role of the object is quite different. In contrast to its absence in the other cases, experience is

something by which the object is (apparently) made present, and which without the (apparent) presence of the object could not take place. Furthermore, we take this presentational aspect of perception to be intimately related to the distinctive phenomenology of perception: perception is experientially as it is because of the apparent presence of the empirical features of things in the experience itself. So whereas the intentionalist claims that in perceptual experience things are conceived of, albeit in a uniquely sensible way, our ordinary assumption is that they are present in experience in a way indistinguishable from that in which we naively think them to be present in the external world. If this were not so, experience could not be the vehicle by which we came to grasp what objects are like in themselves, as we, pre-philosophically, think we do.

There is a response to this presentation-based argument, however. A mental image—unless it has the full force of an hallucination—is not the realisation of a sensible property. It seems to be half-way between a thought and a sensing, and its contents seem, therefore, to be both presentational and intentional. If it is possible in this case, why not in perception? But this response can be turned on its head. The simplest way of understanding the difference between mental images and perceptual experience is that the contents of the former are intentional and of the latter are genuine instantiations of sensible properties. If sense-contents were intentional should it not be possible for them to possess the kinds of vagueness and indefiniteness that characterise images; for the possibility of indeterminacy is the usual contribution of intentionality to thought? But perceptual experience—whether veridical or a veridical-seeming hallucination—remains robustly presentational, and this is what distinguishes it from imaging. The presence of intentionality in the one case and its absence in the other is the only natural explanation of this.

These two points—about psychological vehicles and presentationality—naturally come together. The phenomenological aspect of perceptual experience—the element common to real perception and veridical hallucination—seems more similar to the psychologically real vehicle of thought—the words one says in one's heart—than like the logical object of the thought. The vehicle, we have seen, is not to be explained away as intentionally nonexistent. At the same time, the phenomena seem to be more directly connected to the empirical nature of the object than in other types of mental act. These two facts are naturally reconciled by saying that the object is

present in a quite different manner in perception from other mental acts; not intentionally, but, as far as basic phenomenal properties are concerned, really. The psychologically real ‘vehicle’ of the mental act is the phenomenally realised object, and not a distinct vehicle at all.

This is, of course, a conclusion that the intentionalist cannot accept. But even if they insist on their position, it is a mystery what real purpose is served by deeming phenomenal contents intentional, once the disjunctive account has been abandoned and one accepts an element common to perception and twin hallucination. If I hallucinate a red patch as I look at a white wall, the hallucinatory experience cuts me off from an area of the wall just the size of the hallucinatory patch. In some sense I am aware of red and that veils the world from me. But, whatever the ontological status of the object in the hallucination, it has the same status in perception. If it stands between the subject and the world in hallucination, presumably it does so in perception. This must be clearest in the case of perceptual illusion, when something is looking other than it really is. If some oddity in my eyes or my glasses makes a patch of the wall look red by stimulating a brain process which can figure in either perception or hallucination, then one cannot, without falling into the disjunctive analysis, deny that the red is ontologically on a par in this misperception with its status in the hallucination. Even if the perception is veridical, the status is the same. To call it intentionally nonexistent under these circumstances can be no more than a stipulative decision not to deem it an instance of red, and one that confuses without doing philosophical work.

3 Grammar and intentionality

One way of trying to avoid the conclusion of the previous section, and to give substance to the claim that sense-contents are somehow ontologically soft, is to treat them as, in some way, grammaticolinguistic entities. At first sight, the idea that *phenomena* are grammatical entities seems to be a rather bizarre category error. But we have seen in previous chapters attempts to let language construct experience, and perhaps this is just another instance of that strategy. Anscombe is the philosopher who, most famously, has tried this

approach in her paper 'The intentionality of sensation'. Her argument is lucid and straightforward, if not plausible. She argues

- 1 The question of the ontological status of direct objects—especially the question whether they are 'in language' or 'in the world'—cannot be raised.³
- 2 Intentional objects are a species of direct object.⁴

Therefore

- 3 The question of the ontological status of intentional objects cannot be raised.⁵
- 4 The contents of sense-experience are a species of intentional object.⁶

Therefore

- 5 The question of the ontological status of sense-contents cannot be raised.

Thus the conclusion is that sense-contents are not the sort of things of which it can properly be asked whether they really exist or not. There is also a similar line of argument in play, which puts (2) and (4) together to produce

- 6 The contents of experience are a species of direct object.

Then it adds a new premise:

- 7 The notion of a 'direct object' is a purely grammatical notion.⁷

From which it follows that

- 8 The notion of a sense-content is a purely grammatical notion: 'And "grammatical" is here being used in its ordinary sense'.⁸

From (1) and (7) there follows what is, at first sight, a strange conclusion:

- 9 At least some grammatical notions are not purely linguistic notions.

Although it does not follow from the above premises, Anscombe's argument strongly suggests that grammar in general is not a linguistic phenomenon, but something more mysterious. A certain amount of

mystification concerning the notion of grammar is essential if sense-contents are to be deemed grammatical entities; for I take it that the notion of sense-content is clearly not a linguistic notion, neither as this latter expression would normally be understood, nor as Anscombe understands it, for her idea of a linguistic notion is that it is something which names 'a bit of language';⁹ sense-contents are not thought of as 'bits of language'.

The argument is valid. But premise (1) is false in a way that renders (2) ambiguous; and premise (4), though expressing a truth (namely that I can seem to see something that is not really there), does not provide an analysis of the notion of the contents of experience.

Anscombe's argument for (1) has three steps. First she asks whether direct objects are 'bits of language' or the things for which the bits of language stand. Second, she tries to prove that they are neither. Third, she argues that there are no other 'reasonable candidates' for being their ontological category. We can, therefore, conclude, as in (1), that they do not have an ontological category—there is no sort of existence that they possess. I shall mainly take issue with the second step, though I also have some words of complaint about the arbitrariness of the third.

The natural view is to regard direct objects as parts of sentences—that is, as words or phrases. Anscombe rejects this view, and also the view that they are real objects.

We are given the sentence 'John sent Mary a book'. Anscombe then claims that the four following questions are equivalent:¹⁰

- a What is the direct object of the verb in this sentence?
- b What does the sentence say John sent to Mary?
- c What does the phrase which is the answer to these questions communicate to us; that is, is it being used or mentioned?
- d Is the direct object a bit of language or what the bit of language stands for?

The first three questions are said to have the answer 'a book'; and the fourth to have the answer 'neither, for this is not a proper question'. The reason why *a book* is not a real book is that there is no answer to the question 'which book?', except 'no book', for the example is fictional: but there really is a direct object for the verb. We cannot, however, draw the alternative conclusion that *a book* is a phrase, for if direct objects are merely phrases, and intentional objects are a sub-class of direct objects, then intentional objects are merely phrases. But they

are not, for when I think of or worship Zeus, I do not think of or worship a phrase. In fact there is no need to resort to intentional objects to refute the view that direct objects are phrases, if the four sentences given above are taken as equivalent in giving the direct object, for the answer to (b) is certainly not 'a phrase'.

Contrary to Anscombe's claim, it does not seem that the questions are equivalent. The answer to (a) could be 'the phrase "a book"': the answer to (b), 'it says she sent him a book, that is, some supposedly real book or other': in reply to the further question 'which book?' we can quite happily say 'no book, for the sentence is false in virtue of being a merely fictitious example'; for to say this is not to deny that there was a genuine direct object, that being *the phrase* 'a book'. In answer to (c) we can say that *a book* in answer to (a) mentions a phrase and in answer to (b) uses it. In answer to (d) we say that a direct object is a bit of language, and that question (b) is misplaced, for it does not ask for the direct object, properly speaking. Anscombe says that it does on the basis of how the expression 'direct object' is learnt.

A teacher takes a sentence, say, 'John sent Mary a book' and says: 'What did John send Mary?' Getting the answer 'A book' the teacher says 'That's the direct object'.

This is over-simplified. A teacher at this stage is not concerned to make pupils grasp the distinction between use and mention. Their concern is that the pupil, on seeing or hearing a sentence, should be able to say out loud, write down or put a ring round the correct part of the sentence: if the pupil can recognise what object is said to be, in this example, given, they will pick out the right part of the sentence.

I do not wish to deny that one can use the term 'direct object' so that it no longer refers to a phrase. If one were to ask for the direct object of John's *action*, instead of the direct object of the verb in the sentence, the answer would be, as the answer to (b), a supposed real object. In this case there would not be such a direct object in fact, because the events are only pretend events, though in cases where real actions occurred one could say there were real direct objects of those actions. Thus it would seem that 'direct object' has two uses: one grammatical, in which it refers to phrases or clauses, the other in which it refers to things. The former is the more standard use of the term. When it comes to elucidating the use of intentional verbs,

however, then it is useful to draw on the notion of the direct object of an action, for the root of the peculiarity of these verbs does not lie simply in grammar and the phrases that constitute grammatical objects, but in the nature of certain activities which are strange in that they can be performed upon objects that do not exist. The grammatical object—the phrase—will still exist, otherwise one would not have a sentence: it is what it purports to denote, the object of the *action*, that may be missing.

Anscombe helps to confuse the concept of an *object* of something by talking indiscriminately of objects of verbs and objects of actions. For example, she says

‘Thinking of is a verb for which the topic of the non-existent is full of traps... If I am thinking of Winston Churchill then he is the object of my thought.’¹¹

If ‘object’ has to do univocal service for objects of both actions and verbs, it is not surprising that ‘objects’ dwell in an ontological limbo.

Step (3) of Anscombe’s argument is surely a piece of gerrymandering. Believing that she has shown that direct objects are neither linguistic nor in the realm of things that language is about, she merely refuses to countenance the question ‘what sort of things are direct objects?’ No reason is given for this except that the only answers she is prepared to tolerate have both been dismissed. ‘The direct object of the verb/action’ appears to be a perfectly good referring phrase; if it is to be used in the basic philosophical analysis of sense-experience—that is, if experience is not to be further analysed in non-referential terms, as in adverbialism—then its referential form should be taken seriously. Direct objects could, for example, be located in the realm of abstract objects, with numbers, propositions and universals. If one does not like such entities one should find a way of analysing them out. To be told to attend to the ‘actual use of the term’¹² is not illuminating, for it appears to be used to refer to direct objects. However, as step (2) is fallacious, this issue does not arise.

The collapse of premise (1) is disastrous for the whole of Anscombe’s strategy. First, and most simply, it undermines conclusions (3) and (5). Second, it imports an ambiguity into (2) which prevents the argument to (8). When ‘direct object’ is taken as denoting a phrase in a sentence, then the true version of (2) will be

- 2' The phrase which gives the intentional object in a sentence is a species of direct object phrase.

If we take intentional objects as phrases in this way then (4) becomes palpably false; the contents of sense-experience are not phrases. This prevents (2) and (4) being put together to give the conclusion that the contents of experience are a species of direct object and thus prevents the move from (7) to (8).

If, on the other hand, we take 'direct object' as denoting the thing upon which an action is performed, then (7) ceases to be true in any interesting sense; that is, it is no longer true that direct objects have a 'grammatical' status, either in the sense that they are purely linguistic, or in the sense that they denote things with no ontological standing; for *ex hypothesi* in these cases they denote putative physical entities.

Given that experience cannot be dissolved by the wonderful alchemy of grammar, what is to be made of (4)? Of course, there is a platitudinous interpretation of (4), for it is incontrovertible that at least some experience-reporting statements can be true even though their object expressions fail to refer: thus it can be true that S seems to see a table although there is no table present. As I have already argued, this negative point does not constitute an account of the phenomena, nor, therefore, of what intentional objects of experience positively are. Anscombe's attempt to provide an account by an appeal to the nature of grammar has been shown to fail, and we still have no good sense for 'intentional inexistence' in the case of sensory phenomena.

4 The 'as if interpretation of intentionality

Another possible way of interpreting sensory intentional objects would be in terms of 'as if'. So 'S has a red intentional object' (or 'S seems to see something red') would be analysed as 'it is to S *as if* he were aware of something red'. If this formula is to have any force as an explanation of the intentionality of 'seems' etc., then 'aware' in the analysans must be understood as signifying an extensional relation.

The natural response to any formula of this kind is to ask *how* it is *as if*. 'As if' signifies a similarity, and similarities exist in some respect, so one can ask *in what respect* it is as if one were

extensionally aware. At this point the essentially unilluminating nature of this approach begins to show. The supposed similarity of the intentional and extensional relations consists in the similarities of their objects—it is in being *of red* that they are alike. But it was what it was to be *of red* in the intentional sense that we wished to understand. To say ‘it is like being *of red* extensionally, except there need not be anything red involved’ is to restate what it is for an object to be intentional, not provide an account of how it is possible for a state to be characterised by the nonexistent. This puts pressure to find a way of expressing the similarity that does not make explicit reference to the object. This pushes one down the disastrous path of externalism.

Instead of saying that it is as if one were *aware* of red, one can say it is as if one were *perceiving* red, where perception is taken to be whatever the normal relation is between a perceiver and external physical objects. The idea is that it is *the kind of experience one has when perceiving red*. To do any useful work, this must be understood as indicating a subjective similarity *not* further analysable in terms of the object. Perception is taken as the primitive, and having an experience as of the same object is just to be in a subjective state inscrutably like perception. But if reflection will not reveal the intrinsic nature of the experience in the non-veridical case, it will not reveal it in the perceptual case, which it is like. So the apparent object is reflection-transcendent and external to the mental state. This is the kind of externalism refuted in the previous chapter.

In short, the ‘as if theory is fatally tossed on the horns of a dilemma. Either one explains how it is ‘as if’ by reference to the object, and reintroduces the problem of the ontological status of that object: or one explains the similarity required by ‘as if’ in a way that excludes the object, thereby depriving the experience of its essential content. Experience cannot both do its job as the transparent medium of acquaintance and be opaque and topic-neutral.

These objections apply to the ‘as if’ theory in any context. The situation is aggravated if one accepts the conclusions of Chapter VI. Whilst naive realism is still an option, there is at least a case of extensional awareness for other experiences to be like. But if one accepts the causal argument for the ‘common element’, refuting the disjunctive theory, then there is no real circumstance for it to be ‘as if. This may not be a fatal objection, because one could either (a) allow that there are some genuine extensional relations to sense-data,

though not in those cases that involve the perceptual brain state: or (b) one might argue that it is as if naive realism or the sense-datum theory—that is, as if an extensional theory—were correct, though it is not. The latter seems extremely fishy. It involves essentially characterising theory A—which is supposedly the true theory—as being as if theory B—its rejected rival—were correct. This would create an open season for ontological gerrymandering. One could solve the problem of universals by saying that universals do not exist, it is just as if they did, and leaving it at that.¹³ Strategy (a) is possible but not plausible, given that the avoidance of sense-data is the motivation for the intentionalist strategy. Notwithstanding these uncertainties, the original dilemma is enough to refute the ‘as if theory.

5 Adverbialism and naive realism

The purpose of the adverbial theory of perception is to avoid treating sense-contents as entities and, *a fortiori*, to avoid treating them as entities that stand between the perceiver and the world. This is achieved by not talking about *red sense-data* or *red qualia* or *appearances of red* but instead talking of *sensing redly*. A way of doing something is not an object to which it is done, so the character of appearance is held to reside not in a logical object but in an activity and how that activity is performed. This strategy has been adopted—notably by Roderick Chisholm—as a way of interpreting the ‘inexistence’ of intentional objects.¹⁴ The objects of intentional verbs have their peculiar properties, not by being peculiar entities, but by not being entities at all. They are internal accusatives; facets of the activity, rather than things onto which the activity is directed. So adverbialism can either be considered entirely in its right, or as an interpretation of the intentional theory.

There has been an extensive literature on how to develop the adverbial idiom and about its unnaturalness. ‘Senses redly’ sounds peculiar, but ‘senses redly-squarely’ or ‘red-squarely’ and ‘senses redly-squarely-tablely’ and other variants sound far worse. The objective of the present section is to show that, whether or not adverbialism can be coherently developed, it is of no help to the perceptual realist. Adverbialism faces a problem because, when combined with two other very plausible principles, it entails a conclusion which is both counter-intuitive, and, more relevantly,

more remote from naive or direct realism than is the sense-datum theory.

The first of these principles concerns the relations between physical objects and properties on the one hand, and mental activities on the other. It could hardly be disputed that the only way that physical objects and properties could enter into the content of a mental act would be as the object of such an act; it would make no sense for a physical realist to claim that a physical thing or a property of the sort that are specifically properties of physical things could be a mode of mental activity.¹⁵ So this principle is that types of physical thing and physical properties are not modes of mental activity.

The second principle is that the content of perception is what is ostensibly demonstrable, and, hence, such content is what makes possible the ostensive definition of words appropriately connected with the content. If I know that someone is currently seeming to see a red object, I can teach them the meaning of the word 'red', as the name for the colour that it now seems to them they are seeing. I can do this irrespective of whether the experience in question is a veridical perception, an illusory perception or even an hallucination, so long as I know how it seems to them, for ostension of this sort depends only on the subjective content of the experience, not on further facts about the external world.

The defining claim of adverbialism is that the contents of sense-experience are modes, not objects, of sensory activity. This can be put together with the two principles I have just explained to give the following argument:

- 1 The subjective content of experience is a mode of mental activity, no part of it is a genuine object of experience.
- 2 The subjective content of experience is—or includes—what is ostensibly demonstrable and, hence, that of which the names can be ostensively defined.

Therefore

- 3 What is ostensibly demonstrable and definable is a mode of mental activity.
- 4 Types of physical objects and types of physical properties are not modes of mental activity.

Therefore

- 5 No type of physical object or type of physical property is ostensively demonstrable or definable.

This conclusion is plainly inconsistent with direct realism. It is further removed from naive realism than is any common version of the sense-datum theory. The sense-datum theorist is either a representative realist or a phenomenalist (with which we can classify idealism for present purposes). A representative realist believes that at least some of the properties that are ostensively demonstrable in virtue of being exemplified in sense-data are of the same kind as some of those exemplified in physical objects. And even of those that are not exemplified in bodies, the realist may believe that it is logically possible that they should have been so exemplified. A Lockean, for example, who believes that bodies resemble sense-data only in respect of primary qualities, could hold that bodies might, logically, have resembled them in respect of secondary qualities too. Or someone who held, with Russell of *Problems of Philosophy*, that the physical world resembled the phenomenal only in abstract structure, could hold that it is contingent that it does not resemble it more fully. Even a phenomenalist believes that his phenomenal world resembles in all perceived respects the physical world in which the naive realist believes, so the ostensively demonstrable is not essentially different from the physical either as he, the phenomenalist, conceives the physical, (for, to him, it is just a construction from the phenomenal) or as his principal opponent conceives it.

For the adverbialist the situation seems to be radically different, however. If only modes of sensing are ostensively available, and if the sorts of features bodies possess cannot be modes of mental activities, then it is a category mistake to entertain the possibility of any resemblance between what is ostensively available and properties of bodies. One could as sensibly say that an inert physical body could be proud, intelligent or lazy, as that it could be red or square. Berkeley's maxim that nothing could be like an idea but an idea would be true, and have the modal force it is supposed to have.

This argument, even if sound, does not refute adverbialism, because it is not impossible to accept the conclusion. It involves accepting that we have only a formal conception of the external world, including its spatial properties. We normally rely on visual perception to give us a fuller conception of what space at least might be like, but if the above argument is correct, visual space is a mode of

visual sensing and could have no qualitative, as opposed to structural or formal, resemblance to physical space. The refutation of adverbialism must await p. 180.

6 Adverbial content and sensory object

There is only one strategy I can think of for resisting the argument in the previous section. This is to deny that it is the adverbially construed content of experience which is ostensibly demonstrable and which contains any sensible feature which can be ostensibly defined. This is to deny the second plausible principle, or (2) in the formal argument. As what is ostensibly demonstrable is available in hallucination as well as perception, and as the adverbial account is supposed to apply to all content that is common to both, it is not easy to see how to achieve this objective. The object—that which is ostensibly demonstrable—would have to be thought of as something over and above the content (that is, over and above that which is construed adverbially) whilst the content is held to be all that there is to a veridical-seeming hallucination. The thought would appear to be as follows:

It is by the having of the content that one is aware of the object, but the object is intentionally nonexistent and need be given no ontological status, whereas the content is real.

This formula might even be put forward in a spirit of impatience. The adverbialist could argue that the suggestion—in my second ‘plausible principle’ and (2) in the argument—that what we are ostensibly aware of is the content of perception just blatantly contradicts adverbialism. For that of which we are ostensibly aware is, trivially, an *object* of awareness, and the whole rationale of adverbialism is to deny that the contents of experience are objectual. So the instrumentalism of the formula—the idea that the content is that *by* which we are aware but not that *of* which we are aware—is essential to adverbialism. Otherwise it is merely a rather mysterious thesis about the possibility of talking consistently in a funny way—a verbal version of a Monty Python funny walk.

As it stands, this formulation of adverbialism must be inadequate for it undermines the adverbial theory. If it were correct it would mean that experience cannot be explained by invoking mode of sensing alone, but by mode of sensing plus intentional object. And it

leaves untouched the problem of how something which has no ontological standing at all can characterise a real event. These are problems which adverbialism was meant to solve. Everything in the phenomena was supposed to be in the mode of sensing, so the whole content of a veridical hallucination would have to be in the sensing; and it was taking this idea seriously that generated my argument. The adverbialist now seems to want to say that the content makes us seem to see an object. If this were a causal 'makes' it would solve nothing about the ontological standing of what it causes. But if it constitutes—rather than just causes—the 'seeming to see', we need to know how, for the content is, supposedly, adverbial and hence objectless, whilst we seem to see *something*, which means that that has an object. We are back where we started. Adopting this kind of adverbialism hardly seems better than saying that the sensing projects a non-existent image on a non-existent screen. Some philosophers do, indeed, seem to be prepared to say this; that is, they talk as if it is possible to allow the intentional object to cope with the phenomenology, and let the adverbial account specify the ontology, and say that the intentional object constitutes no ontological problem just because it is only intentional. So Lycan:

I take the view...that phenomenal individuals such as sense-data are intentional inexistents *à la* Brentano and Meinong. It is, after all, no surprise to be told that mental states have intentional objects that do not exist. So why should we not suppose that after-images and other sense-data are intentional objects that do not exist? If they do not exist then—*voilà*—they do not exist; there are in reality no such things. And that is why we can consistently admit that phenomenal-color properties qualify individuals without granting that there exist individuals that are the bearers of phenomenal-color properties.¹⁶

Lycan combines this free-spirited attitude to intentional objects with adverbialism, and he treats adverbial content externalistically, like Tye, whom we shall be considering in the next section (p. 181). I cannot see why, if the above attitude towards intentional objects were admissible, one should not stop there: mental content is constituted by the non-existent, so we need not worry about them even though they are phenomenally real. The fact that Lycan feels obliged to give an account of experience in terms of what actually exists shows that his *via negativa* alone is not adequate. And the

phenomena, too, must be accounted for by what there is, not by what there is not.

The dilemma for the adverbialist is how to make the intentional object be nothing more than the sensing, whilst not letting the sensing be identical with or contain the object: the genie must perform his magic whilst remaining in the lamp. I shall now consider three strategies whereby objects might seem to emerge without ontological cost.

1 It might seem that this problem could be solved by giving a causal analysis of what it is for a mode of sensing to have a certain intentional object. The intentional object of a mode of sensing would be the typical or appropriate cause of that kind of sensing. Such a construal might look suspiciously like the externalist account of content that was dismissed in Chapter V, but it is not necessarily the same. In that case, it was content as a whole that was being analysed causally and externally, with nothing knowable about the internal component of the experience. On the present suggestion, content is internal, but the intentionality of the object is explained externalistically. Nevertheless, this theory does play into the hands of its opponents. Because the intentional object is understood simply as *that which* typically causes the content, and because both (a) the content is what we find in an hallucination and (b) typical causes of content must be contingently related to it, the intentional object will not be directly ostensively definable. What will be ostensively definable will be *what it feels like to see* something red and *what it feels like to see* something square, not the nature of red or square themselves. But this is exactly the point of the argument in the previous section: on the adverbial theory the nature of empirical qualities, as we naively suppose them to be possessed by physical objects, will not be given in experience.

2 A different interpretation of the idea that it is *by* having the content that we are aware of the object whilst the object is nothing over and above the content, would be an error theory. According to this, we mistake the adverbial content for an object. There are various things one might say about this, but it is enough to point out that it does not constitute an objection to my argument. Once one has seen through the error it will become clear that content is not actually the sort of thing that could make physical features ostensively available—it was a mistake to think that it was that sort of thing.

3 One might try a self-correcting error theory. According to this, we interpret the adverbial content objectively, and to that extent make an error about the nature of the content; but the physical properties of common sense really are similar to the way we interpret the content to be, though not to the way it actually is. In this way, the content is genuinely non-objectual, but what is ostensibly demonstrable is objectual, whilst being an artefact of interpretation, not a real phenomenon. The point is that the transition from content to object is a matter of interpretation, not of the nature of the phenomena.

This subtle move is vitiated by ambiguity concerning the role of interpretation. If interpretation restructures the content objectively, then one has not avoided an objectual theory of content, one has merely allowed cognitive activity a role in building up the phenomena: there is no reason why pre-conscious cognitive activity should not go into the construction of sense-data. If, on the other hand, the cognitive activity does not actually re-form the content then after the interpretation there will be nothing objectual to be ostensibly demonstrated. In other words, it cannot be the case both that the interpretation is purely cognitive and, hence, not a stage in the construction of the phenomena, and that what the interpretation gives is something of which one can be ostensibly aware.

7 The refutation of adverbialism

One might be prepared to accept adverbialism, despite its remoteness from common sense. The programme of construing experience adverbially is, however, demonstrably impossible. Frank Jackson seems to me to have proved this.¹⁷

The fundamental rationale of Jackson's argument is that the adverbial theory cannot cope with the structural complexity that the contents of visual experience possess as a result of their spatial—or apparently spatial—properties. Visual experience, whether perceptual or hallucinatory, presents itself as being of a two- or three-dimensional visual field. But a spatially extended field would be a paradigm of an object of awareness. The adverbial theory has to cope with the complexity that follows from the availability of different positions in an extended medium in a way that does not attribute spatial properties to the contents of visual experience. He brings out the problem mainly through what he calls 'the many property problem'.

Suppose that someone is experiencing a red square and a green circle at the same time. They will be sensing redly, squarely, greenly and circularly. But this characterisation is ambiguous between their perceiving a red square and a green circle, or a green square and a red circle. If one were allowed a spatially extended visual field as the medium for these contents, they would present no problem. The difficulty is how to express the appropriate complexity without invoking anything as objectual as an extended field.

At first sight, it might seem easy to solve the problem by an appropriate individuation of the sensings. So someone seeming to see a red square and a green circle will be undergoing two sensings, one redly-squarely and the other greenly-circularly. This is inadequate in at least two ways. First, it is *ad hoc* or even circular. The two sensings are simultaneous and all that distinguishes them are the spatial relations of their contents. It is, at best, unnatural to explain the spatial relations of the contents by reference to the separateness of the acts. It may be worse than unnatural, given that no feature by which the sensings can be distinguished has been given, other than the distinction in content that we are using them to explain. Second, even if this solution worked for the 'many property problem' it would not solve other difficulties that are also grounded in the spatiality of experience. If a separate sensing attaches to each qualitatively different, spatially discrete patch, how are the relations between such patches—for example, that one is to the left of another—to be expressed? If one were to allow that one sensing was to the left of the other, in addition to individuating it by reference to its patch, then they seem to be no different from objects of awareness.

Michael Tye tries to capture the idea that it is a single visual content that is sensed redly and squarely by an indirect or externalist route. Sensing redly and squarely is sensing in the way one does when perceiving a red and square physical object.¹⁸ Thus whether the 'redly' goes with the 'squarely' or with the 'circlely' is not explained by reference to anything internal to the adverbial level, but by reference to the genuinely spatial physical cause with which it is associated. Unfortunately, this places one in the same situation as those externalists discussed in Chapter V. How can one know which external cause is associated with a particular kind of experience, if the experience does not itself possess the wherewithal for enabling one to recognise it, for experience is, in general, the only access we have to the external world? As with some other

forms of externalism, Tye's theory presupposes that we can identify our experiences by correlating them with the situation in the external world, when this would suppose a God's eye view that we could never adopt.

8 The history of adverbialism

The argument of pp. 174–7 showed that, according to adverbialism, that of which we are ostensibly aware is essentially mental. This conclusion might seem to be idealist in tendency—quite contrary to the way adverbialism is nowadays thought of. But this is not surprising, for adverbialism was first put forward as a reply to the argument in G.E.Moore's 'Refutation of idealism'. C.J.Ducasse, who made this reply, was not concerned to defend idealism, only to cast doubt on Moore's argument, which was, indeed, independently dubious.¹⁹

Moore had argued that an act-object analysis of perceptual experience implied the logical independence and, hence, possible separate existence, of the object. It followed that the *esse* of the object of experience was not *percipi*. This was not taken as an argument for naive or direct realism, because the traditional arguments against that theory were still believed. It was taken as showing that the contents of experience were not 'ideas'—that is, not essentially mental—for if they were they could not exist without the mental act. Thus there evolved the idea of sense-data as neither mental nor physical, but as belonging to a neutral sensible category. Once this stage was reached, there seemed to be no objection to these objects existing unsensed, for they were not mental and their *esse* was not *percipi*. So physical space was pictured as containing these peculiar objects, unsensed sensibilia, which were available only to the right perceiver from the right perspective. Ducasse propounded adverbialism specifically in response to Moore's 'Refutation of idealism' and as a way of overthrowing what was thought to be the absurdity of unsensed sensibilia. Ducasse's theory is, therefore, semi-idealistic, in that its purpose is to show that sense-contents cannot exist unsensed.

Descartes compared ideas in the mind to impressions in wax, and adverbialists claim they are being faithful to this picture; for it and they treat contents as modifications of the subject, not something to which the subject is related. At first sight, the theory could as well have been

called 'adjectival' as 'adverbial', which, indeed, the model of the wax suggests. But, for visual phenomena at least, it is not possible to take an adjectival picture literally. The wax could not take on an impression of a shape if it were not itself extended. The Cartesian subject is not extended and even materialists do not usually conceive of sense contents as simple impressions on the body. The adjectival model would have the subject who seems to see a red square qualified by redness and squareness. Adverbialism avoids this, for redlyness and squarelyness are not redness and squareness.

So the original case for adverbialism can be put in the following way.

- 1 Moore has proved that the object in a perceptual act-object state could exist independently of the act.
- 2 The only options consistent with (1) are naive realism, unsensed sensibilia or the rejection of the act-object analysis.
- 3 Naive realism is well known to be false and unsensed sensibilia are unacceptable, so the act-object analysis must go.
- 4 This means treating sense-contents as modes of the observer, which, because the observer is not literally red or square when they seem to see things with those properties, means treating them adverbially.

Ducasse appeared to attack Moore's claim by arguing that there are internal or connate accusatives (that is, objects) as well as external ones; for example, a dance is an internal accusative of dancing.²⁰ Thus far, he could be seen as disagreeing with Moore's claim that all objects are independent of their acts. But he interpreted what it is for an object to be internal in adverbial terms. This fits well with the example of a dance, for a dance is just a mode of *dancing*. So he accepts Moore's conclusion and preserves the act-dependence of contents by treating internal objects as not truly objects, but as modes of activity. Moore, in replying to Ducasse, admitted himself confused.²¹ On the one hand, he admitted to believing that sense-data cannot exist unperceived and, on the other, to finding unconvincing Ducasse's—or any other—explanation of why this should be so. So Moore, in the end, accepted neither adverbialism, nor, in general, the view that sense-data are internal accusatives, nor unsensed sensibilia, but could not explain why sensibilia could not exist unsensed, given that they are *objects* of consciousness.

9 *Esse est percipi* and being an object of consciousness

Can we help Moore out of his problem? We might begin by asking what is the foundation of the problem; why should there be a conflict between objectuality and mind-dependence?

The feeling that there is no real difficulty at all is strengthened by further investigation into Moore's refutation of idealism. The idealists Moore was attacking rested their case on an argument rather like that in section 23 of Berkeley's *Principles*. They argued that something can only be known in so far as we think of it—that is, in so far as it is in our minds—so we cannot conceive of anything external to mind, at any rate, not such that we can have any knowledge or conception of it. The very fact that thought is mental ensures that its objects are mental. Moore's response is that the act—object structure of mental states creates a logical gap between the mental act and its object that leaves open the possibility of the independence of the object: the very fact that everything we think of is an object of thought does not entail that it is a mind-dependent entity. Seeing Moore's argument in this way suggests that there is a fallacy in concluding that objects of thought (or consciousness) *cannot* be mind-dependent. The negation of 'all objects of thought must be mind-dependent' is not 'no objects of thought can be mind-independent', but 'some (which could include all) objects of thought may be mind-independent', which is enough to refute the argument for idealism in question.

There is, nevertheless, the thought that if some objects of thought can be independent of the mind we are owed an account of why some cannot. Perhaps being an *object* of thought creates a presupposition of independence, and an explanation is needed when this is set aside: if a pain or a sense-datum is an object of a mental act and not merely a mode of that act, why cannot it exist independently?

Given that the sense-datum theory involves a dualism, at least of mental and physical states, I do not see that there is any great difficulty in reconciling the act-object analysis with *esse est percipi*. It would be sufficient if there were good reasons why sense-data should be caused by the brain only when the brain is animated by a mind. Exactly what this comes to will vary with different theories of mind. In the case of both the dual aspect and bundle theories, the conscious mind consists only of states of consciousness. In these cases all that is required to make the mind-independence of sense-data impossible is that sense-

data are caused only by living brains which produce them in sufficient quantity for their products to constitute minds. It is now quite trivial that these objects cannot exist without a mind. One might be tempted to argue that such a solution was *ad hoc*, on the grounds that it rested on legislating that only living brains of the right sort could produce sense-data. This objection is misplaced. On any non-physicalist theory of mind, the connection between the brain and the mind is going to be brute; in Feigl's phrase, it will rest on nomological danglers that cannot be integrated with the rest of science. That the production of sense-data should be part of this realm is not at all surprising.

The situation is rather more complicated for other forms of dualism. One might hold that the bundle theory is correct for primitive kinds of consciousness (animals, for example, and perhaps unreflective consciousness in humans), but that full, intelligent consciousness depends on the presence of a true subject with a capacity for thought and self-awareness, which is conscious of the sense-data. On this theory Moore's original argument was correct, but in a harmless way, because sense-data could exist without the act of the real subject, but not outside a mind of some kind.

Another dualist approach could be via a doctrine of embodiment. The self is not extended, but, when united with a body, part of that embodiment is the generation of sense fields and these do possess spatial properties. It is with these that the brain interacts to produce sense-data: the brain itself does not have the capacity to produce mental items, but it need not be thought of as acting directly on a pure unextended ego, but on aspects of the self's structuredness that follow from its embodiment.

All these theories are, in some ways, speculative, but some such theory is required by the non-physicalist and it is no objection to the *esse est percipi* principle that it comes as part of a general non-reductive theory of mind rather than as some more direct analytical truth.

10 Conclusion

The revised causal argument, presented in Chapter VI, together with this chapter, constitute a good argument for sense-data, for they prove the existence of a common element which cannot be construed in an intentional or adverbial way. The argument from illusion in Chapter II was only unsuccessful because of the possibility of appealing to

intentionality and so it, too, is made successful by this chapter.²² As the great philosophers of the early part of this century thought, the Phenomenal Principle is sound, and if I clearly seem to see something red then there is something red of which I am aware. The argument from science is a version of the argument from illusion with a broad scope and it, too, could only be resisted by an appeal to intentionality. It, too, therefore, is successful. Nothing in Chapters VI and VII makes any difference to the other arguments, so the situation with them remains as before.

CHAPTER VIII

The Nature of Sense-data

1 Possible positions on the nature of the common element

The arguments of the previous chapters have, if sound, shown that the disjunctive theory of perception is false and, hence, that there is an element common to perception and hallucination: and that this common element cannot be magicked away by appeal to the ontological elusiveness of intentional objects, or by deploying the adverbial idiom. But this still leaves uncertain the nature of the common element.

I believe that the foregoing arguments establish at least the following minimal position: there is a component in perception which (1) is inconsistent with the truth of a naive or direct realist account of perception; (2) is central to the content of perception, in the way that secondary qualities are; (3) is not physical, in the sense that a typical modern reductive physicalism would require; (4) is not wholly intentional but has, at least, a categorical core. Establishing this much might seem sufficient for at least the moral victory of the sense-datum theory over the kinds of theory usually proposed in its place. But it leaves many questions about the nature of the common element unanswered. These questions concern the relation between this sense-datum or sensational component in perception and the final 'lived experience', which seems to be of a three dimensional, mind-independent, enduring physical world.

Before moving on to investigate the actual nature of the common element, I shall step back and consider all the possible views of it, and then see what is left after the discussion of intentionality in chapter VII.

First, there are two versions of the percept theory. According to the pure percept theory, the common element is entirely intentional, so

that no feature of introspectable content consists of actualised properties. According to weak percept theory, qualities are genuinely actualised in the common element, but other features—such as being of physical objects of various kinds—is only intentionally present. These are, however, part of the phenomenal content and not just creatures of interpretation: this is what makes this still a percept theory.¹

Second, there are two versions of the sense-datum theory. On the weak sense-datum theory, sensible qualities *and depth* are phenomenally real and all other features—such as its seeming to be of some kind of physical object—are matters of interpretation. According to the strong sense-datum theory—which I shall call the maximal position—sensible qualities are phenomenally real, but depth is a matter of interpretation, not a real quality of the common element.

Third, there is the ultra-minimal position. According to this, the common element is not open to introspection. Experience as we consciously and conceptually grasp it does indeed depend on a non-intentional component, but that component is so radically altered by being developed into an intentional structure in the construction of experience that we are not able to introspect its original nature, not even as aspects of the final experience. No scrutable feature of even colour or visual shape, as they confront us consciously, is similar to the nature of the sensational element out of which they are developed.

It is clear that pure percept theory is incompatible with the conclusions of Chapter VII, which showed that at least some phenomenal content had to be non-intentional. The ultra-minimal position combines naturally with pure percept theory, so that there is a non-intentional component in the common element, but everything introspectable is intentional and is somehow generated out of the sensational core by conceptualisation. As Chapter VII was dealing with *the phenomena* it will not be possible to avoid its anti-intentionalist conclusions just by postulating something non-intentional *and subliminal*. As we shall see, however, there are problems with the sense-datum theory that threaten to push one back towards a theory rather like the pure percept theory.

Both versions of the sense-datum theory are still in play, and I shall discuss depth on p. 205. The classic theory is *the strong sense-datum theory*, or *the maximal position*, which is, roughly, that the relation between the sense-datum and the full experience is like that between a

realistic picture—which is, in itself, just a two-dimensional pattern of colours—and the experience one has when one sees it as a representation of whatever it is a picture of. There is an assumption that this latter relation is essentially unproblematic. I shall begin the chapter by investigating problems for the sense-datum theory. This will involve looking at the relation between sense-data, indeterminacy and interpretation (pp. 190–8) and at the phenomenon of depth (p. 205). By the end we should be in the position to develop a positive account and decide the nature of sense-data.

The opponents of sense-data are not yet finished and their hopes centre on the ultra-minimal theory. It is not that ultra-minimalism is necessarily incompatible with the sense-datum theory, indeed it can be combined with either version of the sense-datum theory. On such a synthesis, the proto-phenomenal core is built into sense-data by some process of conceptualisation. I argued in Chapter VII that this kind of interpretative construction was not necessarily inconsistent with orthodox sense-datum theory, because it merely gives a role to pre-conscious conceptualisation in the production of sense-data. However, this attempt to by-pass ultra-minimalism will not work if the final phenomenal product cannot qualify as sense-data, and this indeed will be the case if it could not be said to realise or instantiate any qualities. Thus if, when I introspect a red and square phenomenal object, what I introspect is not actually an instance of redness and squareness, then it is not a sense-datum, but an intentional object. This would, I think, be the only ground for taking ultra-minimalism seriously. Apart from the anti-private language argument, which I have already discussed, there is a collection of arguments challenging the determinacy of sense-contents, and thereby challenging the possibility that they are genuine actualisations of sensible qualities. This gives us an account of sensory intentional objects not discussed in Chapter VII. This is that they are created by interpretation (which means, I think, in this context, just the application of concepts) to a core which, though essentially proto-experiential, exhibits no introspectable qualities itself. This seems a strongly Kantian theory and I find it difficult to decide whether it makes sense to postulate the inscrutably proto-phenomenal. However, if the sense-data theories cannot cope with some of the problems concerning indeterminacy, we could be forced to it.

The ultra-minimal theory will appeal to those who oppose sense-data from a Kantian perspective. I suspect that it might also appeal to

physicalists, though for mistaken reasons. It might seem that ultraminimalism is compatible with holding that the sensational element is physical: for, if we are not aware introspectively of the properties of the sensational element, how can we rule out the possibility that they are physical? Isn't this a version of the theory that our knowledge of our sensations is topic-neutral? This, however, is a mistake. No one is, at this stage, suggesting that the introspectable phenomenal object—that is, the constructed intentional object—is physical, and the nature of sensation has to be such that interpretation of it produces (or 'constructs') the phenomenal object. It, therefore, possesses properties that are essentially proto-phenomenal; and conceptualisation cannot turn brain states into apparent red houses and green trees.

2 Arguments for phenomenal indeterminacy

The intransitivity of exact phenomenal similarity

The problem this constitutes for a sense-datum theory is well expressed by David Armstrong:

'Exact similarity in a particular respect' is necessarily a transitive relation. Now suppose that we have three samples of cloth, A, B and C, which are exactly alike, except that they slightly differ in colour. Suppose further, however, that A and B are *perceptually* completely indistinguishable in respect of colour, and B and C are *perceptually* completely indistinguishable in respect of colour. Suppose, however, that A and C can be perceptually distinguished from each other in this respect.

Now consider the situation if we hold a 'sensory item' view of perception. If the pieces of cloth A and B are perceptually indistinguishable in colour, it will seem to follow that the two sensory items A1 and B1 that we have when we look at the two pieces *actually are identical in colour...* In the same way B1 and C1 will be sensory items that are identical in colour. Yet, by hypothesis, sensory items A1 and C1 are not identical in colour!²

There are two approaches one can take to this. The first is to suggest that perhaps B looks slightly different when it is seen in conjunction with A, from what it looks when it is seen with C. That is, though

one first compares sensory item A1 with sensory item B1, when one looks at A and B, when one looks at C and B, one compares C1 with B2.

This is not at all the artificial escape contrivance it may seem at first sight. Any object, including B, varies in its colour appearance in accordance with its environment—the background or accompanying colours with which it is seen. This does seem to be what is happening here. The presence of A and C influence how we see B—and no doubt the presence of B influences how we see A and C. If the paradox proves anything it is the truth of this analysis. It is, of course, not immediately possible to compare our two views of B—i.e. B1 and B2. But we can perfectly well compare them second hand. B1 was like A1 and B2 was like C1 and A1 was not like C1, therefore B1 was not like B2. As there is no nonsense, and even no empirical implausibility in this view, why not explain the paradox in this way? Once one realises that the same object can have various sense-data, the paradox vanishes.

It might be argued that this escape can work only if (A and B) and (B and C) are compared separately, and not all three together, for this would involve B appearing in two different ways at once, which is impossible. This is true and does represent a limitation on the events covered by my analysis, but it is a limitation shared by the original phenomenon: B could not look like A and like C while A and C were dissimilar, in one ‘act of looking’. Even if the three objects were side by side, one must perform separate acts of looking, attending or noticing, for the phenomenon to occur, otherwise one would have to see C as looking like B, but not like A, although A was like B at the same instant, which involves the contradiction that C both looks like, and does not look like, B.

The second approach involves employing the distinction between sensation and judgement in a way often thought impermissible for the sense-datum theorist. We can say that the judgements about the exact nature of the sense-data are not perfectly accurate; there is a vagueness in our judgements and not an indeterminacy in the sense-data, we just cannot tell what is exactly like what. The usual response to this argument is that it undermines the whole purpose of sense-data, which is to reify how things look, and, therefore, to be exactly how things look: if sense-data can look other than they are why cannot physical objects, and if physical objects can, then we have no reason for postulating sense-data to explain objects’ looking other

than the way they really are.³ This response is too hurried. First, the main argument for sense-data that we used was the causal argument, not the argument from illusion. Second, as we shall see when we discuss 'speckled hen' cases, it is neither necessary nor possible to insist that all *judgements* about sense-data are incorrigible, even if one accepts the argument from illusion. It is a matter of whether variations in perceptual content, or errors in perceptual judgement, (which together encompass the phenomena called 'illusion') can be ascribed to the limitations of our judgemental faculties, or whether they must be ascribed to a variation in phenomenal content which is not a creature of judgement. Perhaps inadequate judgements about similarities can be ascribed to our judgemental limitations, as judgements about the number of speckles certainly can, but the changes that occur if the light alters, or the perspective, or if I take off my glasses, are not changes in conceptual abilities, but in that on the basis of which I exercise those abilities in making my judgements. The argument from illusion would be threatened only if it were plausible to explain what happens when I take my glasses off by saying that it *directly* affects my ability to apply concepts. (Remember, this discussion is not taking place in the context of a theory which says that perception is nothing other than the ability to make judgements; this has already been refuted.⁴ The existence of an irreducibly experiential element is not in dispute. My position is that it is that element that changes in the kinds of cases classed under 'illusion' and not merely our capacities to judge that change. Indeed, there is a remarkable tendency to constancy in that respect, for we are often not deceived by 'illusions'.)

The speckled hen

There is another well-known class of supposedly indeterminate phenomena, represented by the speckled hen. On this one can see a considerable number of speckles, but one is

unable to see exactly how many speckles it has. The hen has a definite number of speckles, but the perception is a perception of an indeterminate number of speckles.⁵

Such indeterminacy is in most perception if not all.

For instance, when I see or feel that one object is larger than another, I do not perceive how much larger the first object is.⁶

It is not correct to claim that a sensory item is indeterminate on the basis that one is unable to articulate all the facts about its intrinsic structure while one is perceiving it. The fact that someone looking at a group of people cannot say how many there are does not show that they did not see them all, in a proper, experiential, non-subliminal sense. Whether one is a sense-datum theorist or a direct realist, one will need a distinction between being aware of something and noticing or excogitating all the comparative and relational facts about the things of which one is aware. Research in perception and counting has, I believe, shown that there is a ceiling to the number of objects that one can estimate at a glance. This should not be taken as showing that, extensionally, one cannot be aware of more than this number of objects.

This distinction between awareness and cognitive articulation will be essential in discussing Dennett's problem cases.

Dennett's long-haired woman

Daniel Dennett raises some interesting cases which he believes show that there is no fact of the matter concerning what sense-contents there are, and, hence, that they cannot be reified. His kind of indeterminacy is more radical than those we have been considering. It is not merely that he thinks that the phenomena are indeterminate and, therefore, not actual instances of qualities, but that there is no fact of the matter about when the phenomena occur. The various interpretations we make 'of it' are all that there is; and there is nothing experiential (or proto-experiential, in the non-physicalist sense required by the ultra-minimal theory) that we are interpreting.

The first case is *the long-haired woman*:

Suppose you are standing on the corner and a long-haired woman dashes by. About one second *after* this, a subterranean memory of some earlier woman—a short-haired woman with eyeglasses—contaminates the memory of what you have just seen: when asked a minute later for details of the woman you just saw, you report, sincerely but erroneously, her eyeglasses ... [W]e are inclined to say that your original *visual* experience, as opposed to the memory of it seconds later, was *not* of a woman wearing

glasses. But as a result of the subsequent memory contaminations, it seems to you exactly as if at the first moment you saw her, you were struck by her eyeglasses.⁷

The interpretation built into this account Dennett calls an 'Orwellian revision':

[T]here was a fleeting instant, before the memory contamination took place, where it *didn't* seem to you she had glasses. For that brief moment, the *reality* of your conscious experience was a long-haired woman *without* eyeglasses, but this historical fact has become inert; it has left no trace, thanks to the contamination of memory that came one second after you glimpsed her.⁸

There is, however, another interpretation of these events, which, because it involves interfering with the evidence right from the start, he calls 'Stalinesque revision'

Your subterranean earlier memories of that woman with the eyeglasses could just as easily have contaminated your experience *on the upward path*, in the process of information that occurs 'prior to consciousness', so that you actually *hallucinated* the eyeglasses from the very beginning of your experience. In that case, your obsessive memory of the earlier woman with glasses would be playing a Stalinesque trick on you, creating a show trial in experience, which you then accurately recall at later times, thanks to the record in your memory.⁹

Dennett thinks that if our intuitive notion of *a conscious datum* were sound, then there would have to be a difference between these two interpretations, but there is not:

There was a time window that began when the long-haired woman dashed by, exciting your retinas, and ended when you expressed...your eventual conviction that she was wearing glasses. At some time during this interval, the content *wearing glasses* was spuriously added to the content *long-haired woman*. We may assume...that there was a brief time when the content *long-haired woman* had already been discriminated in the brain but *before* the content *wearing glasses* had been erroneously 'bound' to it. Indeed, it would be plausible to suppose that this discrimination of a long-haired woman was

what triggered the memory of the earlier woman with the glasses. What we would not know, however, is whether this spurious binding was 'before or after the fact'—the presumed fact of 'actual conscious experience'. Were you first conscious of a long-haired woman without glasses and then conscious of a long-haired woman with glasses, a subsequent consciousness that wiped out the memory of the earlier experience, or was the first instant of conscious experience already spuriously tinged with eyeglasses?

There is no fact of the matter here because

Here the distinction between perceptual revisions and memory revisions that works so crisply at other scales is no longer guaranteed to make sense.¹⁰

As Dennett presents this case, it is possible simply to deny his conclusion and claim that there must have been a fact of the matter, though we are not in a position to tell what. This response is not, however, so clearly available in the case we shall consider next and, therefore, I want to consider another strategy.

Dennett takes his opponents to be saying that there is a particular moment at which an event is made conscious in the Cartesian theatre; some events lead up to this and others follow after it, but it itself is clear and discrete. In fact, Dennett's conception of his opponents' idea of consciousness makes it very similar to the *maximal conception* I introduced above. His assumption appears to be that one has a choice between this maximal conception and his own 'multiple drafts' theory of consciousness, according to which consciousness is *nothing but* layers of processing the 'cash value' of which is simply their behavioural upshot. What he does not appear to allow is that there could be genuine phenomenology, in a traditional sense, but in which the phenomena are greatly affected by the kinds of conceptual activity associated with them. At the one end there may be a kind of bare awareness that is virtually subliminal, at the other a fully self-conscious experience, the fully self-consciousness being the product of the conceptual activity similar to that Dennett has in mind, though not conceived of wholly in his physicalistic manner. Dennett must be thinking that once one has abandoned the maximal picture and allowed that consciousness comes in degrees, then one has no good reason to hang on to the traditional notion at all. This, however, is manifestly

false. Hard cases make bad law and the force of the central cases should not be lost. If one thinks that something over and above a functionalist account of consciousness is necessary to understand the case when one is looking at an unproblematic and stable red patch against a green background (for example), then one must find a way of integrating the phenomenal element into the hard cases. Dennett, however, denies the obvious cases, and we shall look at that on p. 198. Next, we shall consider another hard case.

Kolers' phi phenomenon

Dennett says that this is 'a relatively simple phenomenon that defies explanation by the old theory'—that is, the old theory of consciousness which treats it as consisting in real events in a sort of Cartesian theatre.

[I]f two or more small spots separated by as much as 4 degrees of visual angle are briefly lit in rapid succession, a single spot will seem to move back and forth... The philosopher Nelson Goodman had asked Kolers whether the phi phenomenon persisted if the two illuminated spots were different in color, and if so, what happened to the color of 'the' spot as 'it' moved? The answer, when Kolers and Grunau performed the experiments, was unexpected:...the first spot seemed to begin moving and then change color abruptly *in the middle of its illusory passage* toward the second location. Goodman wondered: 'How are we able...to fill in the spot at the intervening place-times along a path running from the first to the second flash *before that second flash occurs?*'¹¹

Dennett asks whether the colour of the phenomenal spot really changed part-way across, *even though the subject did not know at that stage what colour it was going to be*, or whether the end state is read back into the process to make better sense of it, misrepresenting what was actually in consciousness at the time. He says there is no difference between these for there is no fact about what happened at the time over and above the story that the subject's brain confabulated for him immediately afterwards. If there really were actual contents one or other story would have to be true.

Kolers' phi phenomenon is a particular case of a general phenomenon. The general form of the situation is that, in reality, things

change discontinuously, but in appearance the change is continuous. This phenomenon applies even though the end state is not predictable—say from past cases—things might have gone in more than one way. So the transitional states are filled in in experience after the end state has actually occurred, not by anticipation.

The simplest way of dealing with this phenomenon would be to say that there is such a lag between the physical event and the completed conscious experience that the latter only starts after the former has finished. So the subject becomes aware of the first state only after the last state has occurred. As the lag cannot be very large, (given the speed of light and of nerve impulses) this would explain why this phenomenon only occurs for fairly swift changes—those that take place within what is sometimes called ‘the specious present’. This convenient answer, however, does not seem to be correct. If asked to press a button at the onset of the first event, (that is, the illumination of the red light) the subject starts to press well before the whole cycle is over. This strongly suggests that they start to experience the cycle before it is over. To avoid this, it has been suggested that the button-pressing response is unconscious. But if we adopt the suggestion that consciousness is layered and the phenomena take time to be fully digested—that is, represented in their most interpreted and introspectable form—then the response could be a reaction to the most primitive kind of consciousness and the fully digested version might take place after the actual cycle had concluded. Notice that if the internal time-structure of revisionary experiences like this is not to be distorted, the time it takes fully to process the revised latter part of the experience, when the extra passages are being inserted, would have not to be measurably longer than the time it takes to properly digest the uncorrected initial stages. But as our experience is replete with this kind of phenomena it would not be surprising if it were a normal part of the processing, and so caused no delay.

We can now go back to the woman with long hair. We now have good reason for saying that there must have been an original experience of a fairly developed sort in which she either did or did not have glasses. For the misremembering in this case was one minute later than the woman’s passing by, and a full range of normal-within-the-specious-present processing would have already taken place before the memory: and this processing either should, or should not, have involved some revision.

Dennett presents us with a false dichotomy. He talks as if *either* there is a definite timeable conscious event in a Cartesian theatre, *or* consciousness is no more than the redrafting of accounts of what has been happening in the external world, with no real conscious event. Common sense suggests there is a real phenomenal core which gets conceptualised and interpreted in a variety of ways. The core is needed because without it it is unclear why the various drafts should come out as *experiences*. Dennett's 'multiple drafts' do not consist in internal monologues, nor in any real intellectual occurrences. They are simply layers of 'processing'; that is, physical processes that can be given an intentional interpretation, like the operations of a computer. This entirely loses anything genuinely mental, let alone anything conscious.¹² Common sense does not reject the idea that there are layers of 'processing'—that is, interpretation and degrees of recognition and assimilation—but it does deny both that there is nothing experiential being interpreted etc., and that the process of interpretation is mental only because a certain kind of exegesis can be applied to it. Within that brief or 'specious' present that consciousness can take in as a whole, the experience moves through degrees of conceptual assimilation: without any of this it is hardly a conscious experience *of the subject* for they have not assimilated it.¹³

3 Dennett and the central case: the knowledge argument

I said on p. 196 that difficult cases, like those that Dennett presents, cannot be used to overthrow what is obvious about the straightforward cases. Dennett, of course, denies that there is anything obvious in those cases—he denies, that is, that there is a phenomenal core—so it is relevant to look at his direct arguments for, as I see it, denying the obvious. If he cannot explain the way there seems to be a clear phenomenal core in the central cases, then his use of hard cases to 'deconstruct' our ordinary notion of consciousness must be deemed inadequate.

Dennett's philosophical method could be described as *the Jericho method*. He believes that marching around a philosophical problem often enough, proclaiming what are, plausibly, relevant scientific truths, the problem will dissolve before our eyes. In so far as he is inviting us to adopt a new way of looking at things, this method is quite appropriate. It does mean, however, that moments of direct

philosophical argument are rare, and are to be cherished when found. Not everything can be done by the indirect method, however, and the knowledge argument for the existence of *qualia* condenses the issue sufficiently precisely to force Dennett to enter into a direct argument. If he fails here then his claim that there is no fact of the matter about consciousness will not be plausible and, in view of our treatment of the 'hard' cases, we will be free to give a traditional account of those facts.

Dennett considers the argument in Jackson's famous 'what Mary didn't know' form.¹⁴ Mary knows everything that a completed physical science could tell her about the physical processes involved in visual perception, including colour perception, but has never been allowed to perceive colour, only black and white. If she is finally allowed to perceive colours she will discover something she previously did not know, namely what colour and seeing colour are like. As she previously knew everything physical and relevant, the new knowledge must relate to something non-physical.

Dennett's response is that if we take seriously the premise that Mary knew *everything* there was to know about the physical processes, then the conclusion that she would gain new knowledge on being allowed to see colour does not follow. He continues the story as follows:

And so, one day, Mary's captors decided it was time for her to see colors. As a trick they prepared a bright blue banana to present as her first color experience ever. Mary took one look at it and said 'Hey! You tried to trick me! Bananas are yellow but this one is blue!' Her captors were dumbfounded. How did she do it? 'Simple,' she replied. 'You have to remember that I know everything—absolutely everything—that could ever be known about the causes and effects of color vision. So of course before you brought the banana in, I had already written down, in exquisite detail exactly what physical impression a yellow object or a blue object...would make on my nervous system. So I already knew exactly what thoughts I would have (because, after all, the 'mere disposition' to think about this or that is not one of your famous qualia is it?). I was not in the slightest surprised by my experience of blue... I realize that it is *hard for you to imagine* that I could know so much about my reactive dispositions that the way blue affected me came as no surprise.

Of course it's hard for you to imagine. It's hard for anyone to imagine the consequences of someone knowing absolutely everything about anything!'¹⁵

The message of Dennett's story is that Mary could know what a colour was like because she could work out, on the basis of her physical information, what thoughts she would have given a *physical* colour stimulus—that is, given a certain input of light-waves. This is possible because the *qualia* argument does not deny that a physicalist account of thought could be correct, and so from the physical effects of the light on the eyes she could follow up its consequences in the brain, including in those aspects that constitute her thinking. But if one understands the thought 'that's blue'—which she can because it is just a functional state—then one knows what blue is like.

The first objection to Dennett's argument is that it contains a confusion, resting on an ambiguity. He talks as if Mary could, on the basis of her physical knowledge, recognise the stimulus just by looking at it ('Mary took one look...'). But she will not know by looking at the banana what *physical* colour impression—that is, what input of light-waves—it is giving her unless she already has some way of relating *qualia* and physical states: a complete knowledge of science does not tell one what is going on in one's senses at a given time. It might be argued that she knows what the physical stimulus is as part of knowing 'all the physical information', but here the ambiguity enters. This expression should not be taken to mean that she knows every particular physical thing that is going on, only that she knows all the relevant physical science. It is the general scientific knowledge, the anti-materialist says, that ought to be enough to give knowledge of the nature of experience, if physicalism is correct, and this is consistent with not being omniscient about particular physical goings-on, such as what is happening to the rods and cones in one's own eyes at a given moment. She may have 'written down in exquisite detail, exactly what effect a yellow object...would have' on her nervous system, but she could not tell by looking at an object whether it was having *that* effect, so would not know that *that* look was the yellow—or the blue—one. Dennett's implication that Mary could, through physical knowledge, acquire the ability for direct recognition of colours is mistaken.¹⁶

Nevertheless, there is an argument in the vicinity, and it is very close to Shoemaker's argument that if there are *qualia* then functionalism

accommodates them.¹⁷ Dennett is arguing that from knowing the physical stimulus Mary could work out the physical thought, and, as physicalism is not disputed about thought, therefore she knows the thought-proper, and so understands it and so knows what is involved in knowing its contents, which, in this case, means knowing the nature of the phenomenal colour of which this thought is the recognition. No role is played in this argument by making Mary the subject of the experience. Even whilst deprived of colour experience herself she could work out the functionally defined thoughts that physical stimuli would give others and should, for the same reasons, be able to grasp their content.

There is a truth hidden here but it is not the one Dennett thinks it is. A thought, on his account, is a kind of functional state. It is, roughly, a disposition towards a verbal response, plus other bits of sophisticated behaviour. Knowledge of how someone is disposed to react, verbally or otherwise, does not tell you what it is like to possess a mental state, if there is such a thing as what it is like, which it is *ex hypothesi* there is, otherwise the question is begged. All Mary could know is *what one would say and how one would react* to a certain colour stimulus. So when she sees the banana, if she also knows the nature of the physical stimulus, she will be able to work out (not know spontaneously) that this is the sort of stimulus that prompts 'that's blue', but it will be a revelation that the sort of physical stimulus that she knew was called 'blue' and led to 'blue appropriate' behaviour *looked like that*, and this phenomenal fact is what she comes to know. It is, therefore, not true that she was 'not in the slightest surprised' by the nature of the experience.

It follows that the functional account of the demonstrative thought 'that's blue' does not capture its full content, for Mary can understand the functionally defined recognitional thought without grasping the nature of the phenomenon recognised. The situation is quite different if one is allowed a non-reductive account of thought. Then, provided that the meaning of colour names is essentially ostensive, knowing what blue was like would be part of understanding 'that's blue'.

The important truth hidden in Dennett's (and Shoemaker's) argument is that it is not possible to combine a reductively physicalist account of thought with a non-reductive account of experience, because recognitional thoughts essentially involve the *qualia*. If the functionally defined thought is really all there is to thinking about blue and, hence, to knowing what blue is like, then the functionally

defined thought should be adequate to the cognitive state of recognising blue and the *qualia* becomes irrelevant, like the beetle in the box. Once it has been demonstrated that the *qualia* is not irrelevant, because coming to discover what it is like is a new piece of knowledge, then one is obliged to adopt an account of thought that can assimilate this fact. Those who like Jackson (or Ayer and many traditional empiricists) think that physicalism can be correct for everything but *qualia* are in an inconsistent position.¹⁸ The ‘knowledge argument’ should not be cast in the form ‘physicalism can work for all other mental states but not for *qualia*’, but in the form ‘even if it might look as if functionalism will work for less clearly introspectable states, such as thoughts, Mary’s case shows that it will not work for *qualia*, and we can see from this that it does not work for thought—at least, a certain category of thought, namely those involved in simple recognition—either’.

4 The situation so far

I have considered two different kinds of arguments. My responses to the ‘intransitivity’ and ‘speckled hen’ arguments constituted a defence of the idea that experience could be determinate at a qualitative level whilst being indefinite at a judgemental level. This qualitative determinacy could obtain for ‘fully digested’ experience. My response to Dennett, on the other hand, was consistent with experience being intentional, provided that there is a special sensory kind of intentionality—that is, provided intentionality is not part and parcel of reducing perception to some purely cognitive process. However, there was nothing in what I said against Dennett that was inconsistent with the presence of determinate qualities, and, hence, with instantiation of those qualities in experience. Given the arguments presented in Chapter VII for saying that intentionality is at odds with the presentational nature of sense-experience, the instantiation-of-qualities position still seems the most rational. What remains is to see what instantiation theory to adopt.

5 The maximal position and the blind spot

The simplest sense-datum theory is what I have called ‘the maximal position’, which treats the visual field as being like a coloured photograph; two dimensional in fact but naturally and inevitably

interpreted as three dimensional. It isn't *quite* like a photograph, of course: it gets intrinsically blurred as one moves from the centre and fades away, rather than having a definite edge, for its edge is the limit of experience and not something that we can experience. Nevertheless, these dissimilarities do not seem to destroy the analogy.

I want to consider two problems with this theory, one specific and the other more general. The specific difficulty concerns the nature of the blind spot, and the more general one concerns depth.

If one were aware of the visual field in the way that one is aware of a photograph and if the blind spot were an area of the visual field in which one was aware of nothing, then one might expect to be aware of a blank, black or fuzzy area of the visual field. In fact one is aware of no such thing. But, it could be argued, on the maximal view, this is what ought to happen. The explanation often given for why it does not is that the brain 'fills in' the gap. There are, in fact, two different ways that the brain may be thought of as doing this. On one it generates some extra phenomenal material to fill in the gap, based on what surrounds it: so the picture is completed by painting in the gap. The other option is that it, so to speak, stretches what is already there so as to fill the gap. There should be an empirical difference between these options. If the gap comes in the middle of a complex pattern, then if the latter account is correct, the pattern should be deformed. It should be as if one had cut a hole in some patterned material and then stretched what remained to cover it, or like looking at something through glass with a kink or crack that refracts some of the scene away. It is, however, very difficult to tell exactly which of these things is happening because the blind spot is towards the edge of the visual field where everything is fairly blurred anyway.

In addition to the difficulty of viewing the edge of the visual field clearly, there is uncertainty about the relation between the act of awareness and its content. One might take the phenomenon as showing that the blind spot is a feature in awareness itself, not in, or not merely in, that of which we are aware. So, if the visual field were like a photograph, the blind spot would consist in the fact that one fails to be aware of some area of it: on that ground one may wish to say that that area is blank, for one may not want such an unsensed area of sense-datum.

This approach is challenging to the sense-datum theorist because it leads one to think that if the blindness is in the act, and the act in this

respect determines the content of the field, so should it determine the contents of the other parts of the visual field, and this suggests adverbialism. The advantage of adverbialism in this context would be that questions about the spatial structure of the contents of experience would lapse: one would be sensing F-ly with respect to one place and G-ly with respect to another and not at all with respect to the blind spot; there would be no need to construct a roughly literal picture of the overall result.

The inadequacy of adverbialism has, however, already been proved. The alternative is to deny that visual space is orthodox two-dimensional space. Suppose that one had a natural blind spot in the centre of one's visual field (*natural* because if it were caused by damage it would certainly be noticeable because the neural structure would be set up to expect impulses from it). The area in which it occurred would not then be normally a blurred one but a clear one. What would this be like?

We have so far two major options. One is that the contents would be extemporised, the other that contents would be drawn together, giving a sort of rift line. Perhaps there could be a mixture of both, with some stretching and some patching. The stretching would distort the rest of the field, as with cloth. But there is the option of deviant geometry. Figure 1 represents three possibilities: (a) represents an option which no one supports, because it blatantly denies the phenomena, namely that the gap is observed as blank space; (b) represents the option that the missing part is reconstructed; (c) represents the idea that what is actually perceived is stretched to cover the gap, leaving a fault line.




Figure 1 Some possible topologies for visual space

But there is another option, which is a combination of (a) and (c), and which cannot be pictorially represented. On this option, for the subject there is merely a fault line—that is, a discontinuity of content without a gap—but no spatial distortion. Because the subject is only aware of the outer circle they are not aware of any gap—the gap is outside their awareness—so they are aware of no more than a discontinuity of content. But there is no need to represent this as a continuity in a two-dimensional plane, so there is no need to stretch and deform the content to bring it together. The real phenomenal shape of the field does not, therefore, map simply onto physical space. It can be compared to a doughnut, with a hole in the middle, but the hole is not part of the phenomenal space.

Taking this theory seriously would have interesting consequences of a kind that might placate intentionalists and minimalists. It would suggest that purely phenomenal space is difficult to grasp intuitively—though it is close enough to intuition to be qualified by shape and colour—and that when we think of it, it is by conceiving of it as the ‘physical’ space of the manifest image of the world. That spatiality has, in a very modest way, Kantian interpretative properties, is also suggested by our experience of depth.

6 The maximal position and depth

The maximal position states that visual sense-data are two dimensional and that depth is a matter of interpretation. On the other hand, it is an agreed phenomenal fact that the visual field actually *looks* to have depth; or, at least, does not look flat. The phenomenal facts are, therefore, generally taken to contradict the maximal position. This will only be a contradiction if matters of interpretation make no phenomenal difference. That they do not make this kind of difference might be thought to be what is meant by calling them ‘interpretation’. There is supposed to be a contrast between the sensory or phenomenal, on the one hand, and the intellectual—which covers interpretation—on the other. A clear case of this kind of contrast is interpreting a radar screen. Radar operators are not tempted to see the screen or its contents as a three-dimensional representation of an external scene—unlike a television picture—but they can interpret distance from seeing the spots on the screen. As the example of the television screen shows, the situation in the case of pictorial representations is not like radar. The implication is that, with a television picture, one can recapture by

reflection the fact that the experience is of something two dimensional, just as one can easily grasp experientially that a screen or a picture are just two dimensional, even though it is natural to see its contents as three dimensional. And what makes the difference between this two-dimensional experience and the three-dimensional understanding is, in some sense, purely intellectual, as it is for the radar operator, but in the case of the picture this intellectual difference constitutes a real phenomenal difference: it penetrates and structures the experience itself in a way that it does not do with a radar screen.

Accepting this line of argument involves agreeing that depth is genuinely phenomenal. This, in turn, conflicts with Berkeley's powerful point that depth can never be phenomenal because it is like a line turned end-on to the eye—that is, its contents are in principle obscured. The defining feature of space as revealed visually is that it can be occupied by a visible area—that is, by an area of visible colour. If object *a* is spatially distant from object *b* in the primary two dimensions, then the distance between them is represented by the possibility of various colours being visible in the area between them. But if *b* is behind *a*, then any colour in the space between them will not be visible. Given that vision consists of arrays of colour, it is difficult, therefore, to see how depth can be a part of visual experience in any proper sense.

I can only see two possible ways out of this problem. One is to allow that depth is some extra inscrutable phenomenal property, not constructed from colour. This feels *ad hoc* and bogus, but it faces a more serious objection. Depth is supposed to be a further spatial dimension; so it must be the same sort of thing as the other two dimensions. It should not, therefore, be a *sui generis* feature of experience, but ought naturally to be the same kind of thing as the other dimensions. The other option is to allow that purely cognitive features can alter experience. So the fact that one believes that one object is further away than another makes the scene look different, but there is no adequate way of characterising the difference it makes except in terms of the judgement that it is further away. It is plausible to argue that experience is full of features of this kind. On a hot day, for example, the water looks inviting. This is not a matter of seeing the water and having a certain feeling about it, for the invitingness of the water does not seem to be in my breast but in the water. But nor is it a matter of the water's possessing some difference of visual features—a change in brightness or chroma, for example: my

prepositional attitudes towards the water enter into the overall nature of my experience of it. In the case of depth this could be thought of in more or less reductionist ways. At the purely cognitive end, one might hold that we have a geometrical notion of the third dimension and it is the belief in this that enters into the experience and creates the sense of depth. More reductively, one might think that it is anticipation of how phenomenal objects will move through experience and how one can interact with them that give the feeling of depth. This latter could, but need not, go along with a reductive account of depth itself.

If this line of thought is acceptable, then we could describe depth as experiential, but not bed-rock phenomenal. It is experiential because it is not like the judgement of the radar operator, but enters more intimately into the experience itself. On the other hand, it is not truly phenomenal, because, for Berkeleian reasons, it is not given qualitatively in experience. It seems to be a phenomenological fact that attitudes, beliefs and anticipations can enter into the structure and tone of the basic phenomenal field, without being fundamentally phenomenal themselves.

7 Spatial transmodality: Molyneux's problem

Primary qualities and secondary qualities are often distinguished on the grounds that secondaries are restricted to one sensory modality, but primaries can appear in more. Whether they appear in different senses in a way that is qualitatively exactly similar, or only similar in some more structural manner, is another question. There has been a tendency for empiricists to think that colour and tactile feeling are so radically different that the visual shape that belongs to one and the felt shape that structures the other could not be strictly similar: rather they are features of the two kinds of experience that map onto each other to give a purely structural isomorphism.

The question of whether there can be qualities that are genuinely transmodal in the strong sense is an interesting issue. Molyneux's problem was intended to illuminate the question. Molyneux supposed that someone born blind should acquire sight. Would they be able to recognise, on sight, which shapes were which; that is, would they be able to *see* which shape was the one that *felt* so and so? The idea is that, if they could, then it would be the very same quality in both senses, and, if they could not, then they are not. This

thought-experiment does make the question vivid, but it is also misleading. Molyneux's subject might be able to recognise shapes visually *not* because the data from the different senses are the same, but because we are built—hard-wired—so as to correlate them. And failure of immediate recognition, on the other hand, might mean only that it takes time to orientate oneself if one gains a new sense. So the thought-experiment makes its point, but as a real experiment it would prove nothing.

One could as well express the question in terms of different conceptions of what a quality is. The issue of transmodality is more closely related to the empiricists' nominalism and particularism than it is to their epistemology. Admitting that there could be something in common between a tactile sensation and a visual one looks too much like admitting an abstract idea or a sensible form into one's ontology. Space or extension had better be an abstraction from the way a particular tactile object and a particular visual object—essentially a colour—resemble, rather than a something genuinely present in both. By contrast, an Aristotelian could say that *the very same form* comes from a square object either by sight or by touch. The very same form could be in a phantasm of either kind. (The fact that the Aristotelian phantasm is, for these purposes, little different from a sense-datum shows that what is at stake here is not essentially an issue in the philosophy of perception.) Touch, because we take in only a small spatial array at once, is rather like tunnel vision: an edge or a corner in tactile experience is just like the way it would be in tunnel vision, except that, in the one case, the spatial properties are filled in with tactile qualities, in the other, with colour. The qualitative filler makes no difference to the experience of the shape itself, any more than it would make any difference whether it was filled in with green or with blue. A classic empiricist, on the other hand, ties the nature of an experience essentially to its proper object—that is, to the secondary quality. Different kinds of secondary quality virtually define different realms of being, and it is a form of category mistake to suggest that different kinds of secondary quality possess other directly sensible qualities in common.

The argument for the empiricist position can be put as follows:

- 1 Space as it is presented in any given sensory modality (most clearly vision) is not a purely formal notion, subject to a

complete account in, say, mathematical terms; it also has a qualitative element.

- 2 Any kind of experience that contained an element that was formally like space as it occurs in vision would count as an experience of space even if it failed to share any qualitative component with vision.

Therefore

- 3 What different senses have in common in their presentation of space is formal not qualitative, and the qualitative spatiality in a given sense can be deemed to be particular to that sense.

This argument can be challenged by disputing either of the premises. To dispute the first is to say that the spatial element in perception is entirely captured formally and that the qualitative element is entirely a matter of the relevant secondary quality. This cannot be so, for space cannot be given a complete formal account; it is, therefore, very implausible to suggest that space as it figures in a sense can be so captured. Furthermore, it seems perfectly plain that extension and shape are qualitative components of sight: it is easy to demonstrate ostensively what a red and a green square have in common and this seems not to be something that is purely mathematical.

The second premise can be challenged in either of two ways. First, one might deny that just any qualitative clothing for the structural features of space would give something that we would recognise experientially as space; it requires the right qualities. This seems to me to be implausible. Any kind of experience which underwent qualitative variations that mapped onto the spatial properties of vision would surely present itself as another way of perceiving the same set of objects as were visible. The second strategy would be to claim that whenever the structural properties of space are embodied in a secondary quality then *the very same* qualitative space is presented. So the idea is that space in its qualitative form is a product of abstract structure and secondary quality, with the nature of the abstract structure guaranteeing the sameness of the spatial quality. Abstract structure alone—say the kind of digital representation one might get in a computer or a string of numbers—would not capture the qualitative nature of space, but any genuinely qualitative embodiment

of that structure must generate a uniquely qualitative space, even though the secondary qualities which are the vehicles are entirely different.

It is difficult to judge a theory such as this one. On the one hand, the univocal spatial quality which emerges from the combination of structure and secondary quality seems to offend against Ockham's razor, for one could manage with secondary quality and structure. On the other hand, belief in a common form of space unifies the various fields of experience of the subject.

It is disturbing that it should be so difficult to make a decision between these two accounts. Shape is a prominent feature of both sight and touch and should it not be the easiest thing in the world to introspect our visual and tactile experience and observe whether the way space is presented in them is qualitatively similar?

The only way of explaining this is that the spatiality of a given sense is, so to speak, smothered by our conception of public space; that is, by space as we conceive it to be in our manifest image of the world. In practice, this means that it is smothered by space as we conceive it visually, for that is our dominant sense and forms the basis for our manifest image. Only by relating it to that public space—in our case, to visual space—can we give a determinate form to the spatial aspects of the experience. Hence, if we try to abstract our tactile experiences from the structure of the public (visual) world we cannot find a clear spatial datum. This sounds plausible, but it is not clear how it would apply if someone had a conception of space that did not rely on vision. If, for example, someone were blind, but had, from time to time, experiences of flashes of colour with, as we would say, a spatial structure, (one flash to the left of another, for example) could they fail to see the structuring that we express by spatial language ('*a* to the left of *b*'), even if they did not notice that this was the same as what they call spatial? It is difficult to imagine how there could be any kind of visual experience without this visual—spatial feature being blatant, for such extension is of the essence of colour and, hence, of anything that could make it proper to call it visual. This suggests that even if it did not occur to the blind man with spatially structured visual experiences that their structure was spatial, nevertheless he would have a clear sense of what they were like; so that, if they did happen to fit with the spatial structure of his experience it would be possible to recognise the form that spatiality took in his visual experience. It would never, that is, be like touch,

where it is difficult to tell what spatiality there presented is like. If the Aristotelian were correct, someone who did have two unco-ordinated realms of spatial experience should, on reflection at least, realise that they were both spatial in the same sense and naturally project the contents of the lesser one (assuming one to be dominant) onto the other. It should, therefore, be natural for the blind man with the visual hallucinations to see the flashes as occurring at certain points in his public—that is, essentially tactile—space. Public space here includes the subject's own body, so they need not be tempted to see the flashes as outside them. It would be enough if they saw them as being, for example, one in the left side of their head and the other in the right. This would be analogous to the way our sensations are spatial but subjective.

These difficulties, however, do not prevent us from concluding that the reason why it is difficult to solve Molyneux's problem by introspection is that our tactile experience is so ill-defined, in comparison with our visual experience, that we cannot discern such features in it alone.

What are the consequences for us of allowing the possibility of transmodality for spatial perception? Because the empiricists were so keen on disallowing it, it is tempting to think that it may be inconsistent with the sense-datum theory. This, however, is a mistake. Data from different senses have a common temporal structure that is more than purely formal, and this does not harm their standing as separate data. One might dismiss this by saying that temporality is a property of the experiencing rather than of the data themselves, reflecting Kant's idea that time is the form of inner sense and space of outer sense, that is, of the data themselves. It is true that it lacks even *prima facie* sense to imagine constructing a common experiential time from qualitatively different time data in different modalities, for one needs the time at the heart of one's subjectivity to be carrying out such a process; spatiality is not at the centre of oneself in the same obvious way. But the fact that spatiality is a property of the data does not mean that allowing transmodality harms the sense-datum theory. It does mean that instead of several spatial fields from which a unified field has to be constructed, the subject naturally has one spatial field in which various radically different kinds of data can be located.

8 Conclusion

We can conclude that there is nothing wrong with the strong sensedatum or maximal position, with two small modifications that reflect a grain of truth in intentionalism and ultra-minimalism. First, we must agree that interpretation can make real experiential differences, especially in our experience of spatiality. Second, and closely connected, purely phenomenal space is not properly conceptualisable, at least in the normal ways that we think about space, though, at the centre of the visual field at least, it can be thought of as Euclidean and two dimensional.

CHAPTER IX

Sense-data and the Physical World

1 The options

Given that we have established the existence of sense-data, there now arises the question of their relation to the physical world. There are two generic types of theory that claim to answer this question. According to one group of theories, sense-data are caused by, and, in some sense, represent the physical world. According to the other group, the physical world is no more than the experiences—actual or actual and possible—that conscious subjects can have of it. In this chapter I shall investigate versions of both representative realism and phenomenalism/idealism and defend them both against certain standard objections.¹ I shall suggest that the phenomenalist and idealist theories have some important advantages over representative realism, but this conclusion will be tentative because this book is concerned with the philosophy of perception not, essentially, with the nature of the physical world.

2 Varieties of representative realism

Representationalism can come in three degrees:

- 1 According to the most straightforward and strongest version of representationalism, sense-data possess both primary and secondary qualities and these are systematically related to—in the most straightforward cases, directly resemble—those of the objects they represent.
- 2 A weaker and more common form of the theory holds that sense-data resemble their causes only in respect of primary qualities.
- 3 The weakest theory asserts that not merely is there no

resemblance on the level of secondary qualities, but that the primary qualities exhibit only a structural isomorphism.

The arguments specifically against (1) are the arguments for denying that objects possess secondary qualities intrinsically, and I have already discussed these in detail.² The upshot was that the arguments were powerful but not conclusive. Once one has abandoned naive realism, however, the motivation for holding onto the objectivity of secondary qualities is greatly diminished. I shall not, therefore, discuss this option further.

The classic objection against (2) in particular is Berkeley's claim that the idea that primary qualities could exist without secondary qualities makes no sense.³ Berkeley's reason for this is that one cannot *conceive* of primary qualities without secondary qualities, and, for Berkeley, this only means that one cannot *imagine* or *image* primary qualities on their own. Berkeley's argument, that is, entirely depends on an imagist theory of thought. Given that the imagist theory of thought is totally inadequate, Berkeley's argument collapses. There is, however, a related argument that does not depend on imagism. A world that consisted only of primary qualities would possess only geometrical properties and something like impenetrability or force. This is equivalent to the claim that the world consists of mutually resistant forces arranged in space. There are well-known arguments that this conception of the physical world is inadequate because it attributes nothing categorical to the world except empty space. Assuming that there is force to these arguments, it follows that matter must possess some further qualitative filling beyond the primary qualities. It does not, however, follow that this must be our familiar secondary qualities, only that it must be something qualitative and somewhat analogous to the secondary qualities. It would have to be in principle unknowable, for science only reveals primary qualities and anything revealed by sense which is not a primary (or tertiary in Locke's sense) quality must be a secondary quality. This matter-filling quality, therefore, looks suspiciously *ad hoc*, but the conception of it—given that one is not obliged to be able to imagine what one can conceive—seems coherent and, from a realist viewpoint, probably necessary.

The third version is found in Russell's *Problems of Philosophy*. The rationale is an extension of that which operated in Molyneux's problem.⁴ Just as Molyneux and Locke thought that shape as represented in vision and touch would have to be qualitatively different,

so shape as it occurs in physical reality either cannot be, or we have no reason to think it is, qualitatively similar to the form it takes in any experience. All that we can know about shape or most or all other primary qualities in their physical form is what can be captured in some abstract—probably mathematical—form. There would presumably have to be more to reality than is expressed in the mathematics, for mathematics is purely abstract, but what that more is, we can never know. This makes the physical world very remote, but seems a coherent position.

3 General objections to representative realism

Epistemological

In addition to the objections to particular theories already considered, there are two objections that can be brought against any version. The objection most often employed is the epistemological one.⁵ This is that, once we have conceded that we are directly aware only of sense-data, then we lose any reason for believing in the physical world; we are cut off from it by a ‘veil of perception’ which neither sense nor reason can pierce. This argument draws its appeal from the image of the ‘veil’ but its argumentative force is much more esoteric, requiring strongly sceptical principles of a Humean kind.

The intuitive response to the epistemological objection is that the veil of perception presents no threat, because belief in the physical world is required to provide an explanation of why we should have sense-data, and why those sense-data should seem to be of a physical world. When you watch television you see the screen, not the people portrayed, but it is much the simplest theory to believe that the people are there, at the other end of a causal chain that gives rise to the picture. This is the Lockean position. Hume’s challenge is to ask how such an explanation is supposed to work. It is obviously not possible to argue deductively for the existence of an experiencetranscendent physical world, so any such argument would have to be inductive. But Hume has a barrage of objections to this. First, he holds that inductive argument, though instinctive, has no rational force. Second, induction concerns the prediction of so-far unexperienced events on the basis of ones experienced previously; it does not, that is, break out of the world of experience. To employ it in an argument for the existence of the external world we would have had, at some time in the past, to have

experienced a correlation between experience and the world beyond the veil, and to project this correlation. But the problem is to get to a world beyond the veil in the first place. The natural response to this is to concede that what is wanted is not, in the strict sense, an inductive argument, but an argument to the best explanation. What we want to explain is why there should be experience at all—or, perhaps, just highly ordered experience—and the postulation of an external physical world is the best and simplest explanation of this.

Humean scepticism will have nothing to do with such explanations. First, according to Hume causes explain nothing, because causes do not bring anything about. A causal ‘explanation’ merely brings a kind of order to our understanding, it does not say what *makes* anything happen, for nothing ever *makes* anything happen.⁶ A postulated physical world would not, therefore, explain why experiences should occur by saying what brought them about: but neither would such an explanation bring order to our ideas, for the whole point of it is that it seeks to relate our ideas to things that wholly transcend our minds. But the second point is more crucial. Any search for an explanation presupposes that there is something in need of an explanation—that is, something which is improbable unless explained. In this case, this would be so only if it were improbable that it should be a brute fact, without any explanation, that our experience should be highly ordered. We do, indeed, have a strong intuition that this is so. Our experience is sufficiently highly structured for us to be able to interpret it as experience of a fairly stable physical world. Our natural response to this is to assume that experience has a cause: it doesn’t just happen in this ordered way by accident. In arguing this way from experience we are presuming that the fact that order requires an explanation is an *a priori* principle; we do not *learn a posteriori* that the order in the phenomena has a cause, for we have reason to believe that experience has a cause only if we take its orderedness as being evidence for such an hypothesis. The Humean tradition rejects the application of *a priori* probabilities to empirical questions.⁷ On the other hand, belief in such an application is as difficult to doubt as it has been to justify. Common sense argues that if there are no causal constraints on a succession of events then those events should not exhibit clear and distinctive regularities. We assume that, in the absence of causal influence, all possibilities are equally likely and that no striking regularities should emerge. The sceptical rebuttal of this commonsense assumption is well known. The sceptic argues that a highly ordered sequence is, *a priori*,

no less probable than any other. Taking the example of throws of a die, a succession of ten throws of six is no more improbable than a succession of any other numbers, including apparently random ones; e.g. 2, 4, 1, 4, 6, 3, 3, 2, 5, 1. This is because at each throw each number has a chance of one in six, so a sequence of sixes and the random series each have a probability of $1/6 \times 1/6 \times 1/6 \dots$ for each throw; that is, of $1^{10}/6^{10}$. It follows from this that an apparently random sequence of experiences would have just the same improbability as an ordered sequence. *A priori* probabilities merely state how many possibilities there are and assign a probability to each of one over the total.

Although this argument often fails to carry conviction, it has proved difficult to refute. Indeed most non-sceptical philosophers, especially ones not expert in probability theory, are tempted to follow the strategy of the Scottish minister preaching on the problem of evil ('Brethren, this is a very difficult problem, which we must look squarely in the face, and pass on'). Their intuition that order requires explanation remains undimmed, though they abandon the task of justifying it. I, too, am tempted by such involuntary modesty and what I have to say is extremely tentative.

Perhaps one can solve the problem by switching attention from the probability of particular series to the probability of types of series. What we are interested in is the probability of a series which is so structured that a conscious subject is able to discern manageable recurrent regularities. There are obviously many more series of logically possible experiences that do not fit this requirement than ones that do. What we are interested in is the probability of a series *qua* member of such a class. As there are many more series which do not fit this requirement than ones that do, *a priori* it is improbable that one fitting the requirement would occur by chance.

A standard reply to this is that we can always invent a class for any given series which makes it improbable in this way. Thus supposing the series 2, 4, 1, 4, etc. actually occurred, we could say that it belonged to the class containing it alone, or a class containing it and another equally undistinguished series, and argue that it was *a priori* improbable that a member of such a class should occur. It is essential to my suggested approach that it be true and important that such classes do not pick out a genuine kind of series, but are only classes connected by enumeration. Assuming (as I shall) that there is a difference between genuine kinds, whose members are really similar,

and artificial classes, then classes of the arbitrary sort are just irrelevant to questions concerning whether a particular kind of series is probable or not.

Another possible objection is that there are an infinity of possible ordered series, as well as an infinity of possible non-ordered ones. If this is so then neither is more probable than the other. This argument has received an interesting reply from Armstrong and McCall.⁸ They imagine God running a lottery for places in heaven in which every positive integer is represented on a ticket. On the principle that all denumerable infinities are the same size, an offer of a ticket covering all even numbers would give you no better a chance of winning than one covering all numbers divisible by a million. In case one feels intimidated into thinking that unease at this shows only that we are not used to thinking in terms of infinities, they suggest a choice between a ticket on which one wins for all numbers divisible by three and all divisible by six. As the former includes all the latter and an infinity more, it is impossible to see, intuitively, how the former could fail to be a better bet than the latter. I find this very convincing, and it suggests that infinity does not *affect proportions*, which is all my argument requires. The limiting case would be to be offered either a ticket which wins on every number or one which wins only for evens. As the evens can be mapped one to one with all the positive integers, if all infinities are equal in the relevant way then one stands just as good a chance betting on half the numbers as on all of them. As one could not lose in the latter case and could in the former, this must be false.

A further reservation about my argument might be based on the suspicion that the introduction of kinds of series in preference to individual series as bearers of probability is gerrymandering. Why aren't the probabilities of individual series relevant too? I think the answer to this might lie in the fact that probabilities are comparative and an improbability is only salient if it shows an event to be improbable relative to the other options. All series taken as individuals are equiprobable, so the outcome would not have been less improbable if a different series had occurred. From the perspective of genuine kinds of series this is not so, and on this basis series can be differentiated.

I leave this issue with no confidence that I have advanced it. Even if I have not, one can still fall back on the strong intuition that there must be some explanation of order, though the intuition is not yet

justified. Popper rejects solipsism on the grounds that he is sure that he could not have produced the music of Bach; to establish the principle that manifest order is not self-explanatory we need only add to Popper's modesty a conviction that the music of Bach did not write itself.⁹

Ontological

I call the second set of objections 'ontological' because they concern not the difficulty of justifying a belief in an experience-transcendent world, but problems with giving content to the supposed nature of such a world. Berkeley began this line of thought by arguing that we could make sense, neither of the idea of an unperceived object, nor of an object that possessed primary but not secondary qualities. Both these objections, in their original form, rested on a conflation of what we can conceive and what we can image, and nothing can be done to rescue imagist theories of thought. Nevertheless, the spirit of Berkeley's attack on representationalism can be rescued. Two genuinely Berkeleian thoughts can be used to generate arguments against a transcendental material realm beyond the veil of perception. First, there is the thought that we can give no adequate content to our concept of such matter. Second is the conviction that our conception of the physical world is so tied to the manifest and discoverable structure of the world as we experience it that nothing transcendental and possibly remote from that structure could count as the physical world. These lines of thought are mutually supporting and I shall begin with the first.

Apart from undermining the intuitive legitimacy of realism, the main significance of the refutation of naive realism is that it makes it implausible to impute non-dispositional secondary qualities to matter itself. This leaves us with the problem of trying to decide what properties matter is supposed to possess. Notoriously, the Cartesian idea that it is purely geometrical will not do for it leaves no distinction between matter and empty volumes: a filler for these volumes is required. Equally notoriously, Locke's filler, solidity, will not do the job, for that quality collapses on examination into a composite of the dispositional-cum-relational property of impenetrability, and the secondary quality, hardness. What the physical realist requires is clearly put by Harré:

Solidity is the alleged quality, the possession of which is responsible for the fact that two material things cannot occupy the same place at the same time and is logically connected with impenetrability, the power to resist penetration, in that the possession of the former is supposed to account for the manifestation of the latter.¹⁰

Mackie appears to believe that this is possible. He says that our conception of solidity is

only the indirect and relative notion of it as the supposed or inferred ground of a power which is itself learned from its manifestation.¹¹

We can presume here that Mackie intends 'ground' to signify something stronger than contingent connection, otherwise the grounding of impenetrability would no more explain impenetrability than would colour if all impenetrable objects were coloured. I take it that Mackie is after something of the sort Harré describes. But how can there be a necessary connection between a dispositional and a categorical property? Such a connection can be made trivially by use of a 'bridge concept'. A puncture is by definition a double event consisting first of the breaching of a surface, then of the escape of a contained gas or liquid, but the existence of this concept does not explain the connection of the first event with the second. Similarly, calling something 'solid' if it possessed quality Q in association with impenetrability, would not explain how the former grounded the latter. Mackie seems to be confused about what would constitute the prescribed sort of necessary connection. He says:

Modern physics will not use [solidity]; but electric charge is one feature which has come into physical theory to play a corresponding part, and mass (rest mass) is perhaps another.¹²

Electric charge is akin to concepts such as energy or field and is a dispositional or power concept. Mass, too, is either defined in, for our purposes, a question-begging way as a 'quantity of matter' or simply records how objects behave in their interactions. Now some philosophers have claimed that the physical world consists essentially of relational properties—generally of active powers or fields. Hume's objection to conceiving bodies as volumes of impenetrability is still essentially the objection to such conceptions,

namely that they are vacuous. An object cannot simply be a spatially extended capacity to effect other spatially extended capacities to effect... An ontology of mutual influences is not an ontology at all unless the possessors of the influence possess more substantial features.

It might be objected that the claim that there could not be a categorical property which was logically connected with a power is an *a priori* prejudice, probably inspired by atomistic and empiricist assumptions: we cannot legislate *a priori* about what types of property there may be. But the objection to this sort of property is not based upon such an appeal to intuition.

It is trivially true that there is no property P the possession of which analytically entails the possession of a power Q, P not itself being explicitly a power. As the connection would not be nominal or analytic, it would have to be *de re a posteriori*. In fact two sorts of such connection are sometimes considered. The first is not properly *a posteriori*, but rather lives in the shadow cast by the 'paradox of analysis'. There are some truths which are not analytic in the sense of being trivially obvious or merely verbal, but which are still knowable *a priori* by conceptual analysis. Those physicalists who seek to prepare the way for identifying mental with physical properties by giving the former a topic-neutral analysis see truths about mentality in this way. One way in which a necessary connection might be discovered between some non-power property and a power would be if a philosophical analysis uncovered an unobvious connection between them. I do not see how this method could lead to a solution in the present case, because I do not see how, once the non-power concept was fully analysed, if it revealed a powerentailing element that could be other than contingently associated with its non-power aspect. Anyway, I know of no such analysis. The second and properly *a posteriori* necessary connection which is nowadays discerned between concepts would be some form of *de re* necessity. It is often said, for example, that a *de re* necessity links the property of being hot and that of possessing a high mean kinetic energy of its constitutive molecules. Such necessities, however, seem to be more a matter of verbal definition than is generally allowed. Even if we decide to annex the term 'heat' to phenomena associated with mean kinetic energy and 'mean kinetic energy' to behaviour which includes the normal phenomena of heat, it does not follow that there could not be phenomena in all other respects just like either of them yet lacking

this connection. (Indeed, if Kripke is correct in his explanation of how such identities can be *a posteriori*, there must be such possibilities.)¹³ Heat and mean kinetic energy, therefore, seem to be only 'bridge concepts', like 'puncture', and do not non-trivially explain the connection between the two sets of phenomena which belong paradigmatically to each. The same would apply to any attempt to explain impenetrability in terms of any underlying quality or structure.

The situation is, therefore, that the following three negative propositions concerning the nature of matter can be proved beyond reasonable doubt: (1) that matter does not possess non-dispositional secondary qualities; (2) that matter does not possess any sort of primary quality which could logically ground its basic powers, for example, its impenetrability; (3) that matter cannot consist solely of spatially arranged powers or dispositions. As far as I can see this leaves only one possible conception, namely (4) that matter possesses some unknown quality or qualities, conceived on analogy with a sensible quality such as colour. As we have shown that there cannot be an internal or necessary connection between an intrinsic quality of this sort and causal properties, the causal properties of matter must be dictated by laws or powers which are only contingently connected with the matter: that is, there is nothing about the intrinsic nature of the matter which determines what laws govern its conduct, nor to prevent its being governed by different laws at different times.

We have done all we can to develop the first Berkeleian thought, which was that we can give no adequate content to our concept of matter, and have forced the realist to (4) in their defence. Against that we now bring to bear our second Berkeleian thought, which was that a transcendental world beyond the veil of perception would not be the physical world that we inhabit.

The picture of the physical-cum-empirical world which the present conception of matter, when combined with the refutation of naive realism, requires is a 'two-world' picture. There is the world of transcendental physical reality, which consists of objects possessing the unknown qualitative nature, and there is the common, collective or intersubjective phenomenal 'world', which is the world-as-experienced. The latter world is usually thought of as representing the former by resembling it, at least in structural or abstract ways. One thing which is meant by saying that the phenomenal world represents or resembles the transcendental physical world is that the scientific

laws devised to apply to the former, if correct, also apply (at least approximately) to the latter. Physical science, though developed through our acquaintance with the world-as-experienced, aspires to describe the formal features of the world-in-itself. On the present theory of matter, however, there is no need for there to be such a match between the laws discovered by empirical science and those which govern the world-in-itself. The world-in-itself consists of the qualitative core contingently governed by certain laws. Amongst these laws are those which determine how the world will appear. In this way the transcendental physical world gives rise to the empirical world. But it is logically possible that there might be many different exclusive sets of laws which fix how the world should appear. It follows that, from how it appears, we cannot infer how it works in itself. There is an uncertainty in our conception of how the transcendental world is which parallels that which is generally supposed to exist in our explanation of actions from beliefs and desires. It has been argued that different combinations of imputed beliefs and desires will explain the same action, and, as we have no independent way of fixing either parameter, the explanation is holistic and undecidable. Similarly, our experience is the product of (a) how the world-in-itself really works and (b) the laws transforming that into experience. *Ex hypothesi* we have no direct access to either of these things, for we experience the product, the intersubjective empirical world. Remember that the qualitative core of itself entails no particular laws, for the qualitative core intrinsically possesses no causal properties, not even of solidity. It is as if the transcendental world was in a code which is being interpreted by a machine to the internal working of which we can have no access. The messages that reach us are clear and consistent, but that is quite compatible with the original code being changed daily, so long as the principles on which the interpreting machine works also change. It follows that it is both epistemically and logically (or metaphysically, as some say in such contexts) possible that the transcendental world operates on laws quite different from those which a perfect science of the empirical world would generate. That it is logically (or metaphysically) possible follows from the fact that the intrinsic nature of the world (i.e. the unknown qualitative nature of its matter) entails nothing about the laws of its operation and hence provides no constraints on how it should operate and hence on whether its laws of operation should alter. (It is not reasonable to say that the identity of the

transcendental world changes if its laws change, because the very same quality-bearing objects could come to be governed by different laws: if they are the same objects then it is the same world.) It is, therefore, not merely the case that we do not know whether the transcendental world follows the laws derivable from the empirical world, but that, even if it does, it is logically and metaphysically possible that it might cease to do so, without our being any the wiser.

This creates a problem for the physical realist. The physical realist's conception of the physical world has two components. The first is that the world is mind-independent; the second is that the physical world is what physical science investigates and hence its nomological structure is what physical science approximately uncovers. Both these are necessary truths for the physical realist. It is the former which distinguishes his position as realist, but the second condition is also essential, both (a) because it is a conceptual truth that the physical world is what physical science investigates, if it investigates anything real, and (b) because if mind-independence is severed from being the object of science then 'physicality' comes to mean simply 'non-mentality', and to have no positive content of its own. But the concept of physicality is not that thin. It combines both the idea of being 'out there' with that of being investigable by us in certain ways. A world not accessible to certain sorts of investigation, paradigmatically those exhibited by physical science, is not the physical world. Once naive realism is refuted these two conditions come apart, for the following reasons.

The representative realist holds that the transcendental world is the physical world. It is a necessary condition for the truth of this that the transcendental world have approximately the same nomological structure as a developed physical science would attribute to the physical world. An argument against physical realism might initially proceed:

- 1 As I have shown it is possible that the transcendental world does not realise, even approximately, physical laws.
- 2 It has already been established that realising physical laws, at least approximately, is a necessary condition for being the physical world.

Therefore

- 3 It is possible that the transcendental world be not the physical world.
- 4 Identities are necessary.

Therefore

- 5 The transcendental world is not the physical world.

The necessity of identities is essential to the argument. The representative realist's natural inclination will be to say that the transcendental world is the physical world provided that the transcendental world embodies scientific laws; if it does not then there is no physical world, however things may appear. This could be maintained consistently with the necessity of identities provided that the realisation (or not) of scientific laws were an essential feature of a given transcendental world, so that if the actual transcendental world embodies physical laws, it then qualifies essentially as physical. But the contingent connection between laws and ultimate objects means that it is a contingent feature of a given transcendental world whether or not it realises a given set of laws. Suppose that the transcendental world realised physical laws and thus itself counted as physical, it might cease to realise them. The realist would then have to say that the physical world had ceased to exist, though the transcendental world, which had been the physical world, endured. Wiggins has, I think, shown that identity cannot work in this way, provided that the referring expression in question (in this case 'the transcendental world' and 'the physical world') pick out a substantial and complete object.

The transcendental world could be, then cease to be, the physical world only if being the physical world was a matter of playing a certain role, like being Lord Mayor, not a matter of being a particular thing or substance. But if we adopt this view of being physical, then it follows that what is physical (what 'plays that role') is not essentially physical. First, this is not our normal (especially not our normal realist) conception of the physical world: and second, it concedes to the non-realist that the physical world is a logically secondary product of the action of something not itself physical. Additionally, the representative realist's premiss that there is no physical world unless it is the transcendental world is not plausible. Suppose that initially the transcendental world realised physical laws, but then came not to do so, though the world remained unchanged from the viewpoint of all

conscious beings within the natural or empirical realm. It seems wrong to say that the physical world would have ceased to exist just because the way it was sustained by the transcendental reality had changed. If that change were gradual, there would be no point at which it was natural to say that, unbeknownst to us, the physical world had ceased to exist. If this is accepted, then being the physical world is not just a role played by the transcendental world, because the physical world can continue when the transcendental world ceases to play that role.

In sum, the argument is this: acceptance that if there is a physical world then that is what physical science investigates, and observation of the logical features of identity, together show that what lies behind the veil of perception cannot be the physical world. Naive realism, Berkeleyan style, is correct; the physical world is the world of experience, irrespective of what more distant reality sustains it. And the world of experience is essentially mind-dependent because given the falsehood of the physical realist version of naive realism the immediate objects of experience are mind-dependent. This seems to me to be a fairly powerful argument.

4 Types of phenomenalism

It is useful to divide phenomenalism into three kinds. These are theistic phenomenalism, sceptical phenomenalism and analytical phenomenalism. The first is represented by Berkeley, the second by Hume and the third by most twentieth-century phenomenologists. The third is not exclusive of either of the other two. Indeed, it would be commonly held that it is something to which all phenomenologists are committed, though we shall see reasons to doubt this. The first and second agree that our sense-data come in a way that enables us to interpret them physically but disagree about the explanation of this. To be more accurate, they disagree about whether this fact *requires* an explanation at all. The sceptical phenomenologist accepts the orderedness of experience as a brute datum which neither needs nor is susceptible of explanation. The rationale for and weaknesses in this position were discussed on pp. 215–19. The theistic phenomenologist accepts the *a priori* improbability of unexplained order and thinks, for whatever reason, that we cannot explain it by postulating a mind-independent physical world. It is therefore explained by postulating an immaterial agent.

Analytical phenomenalism is the most difficult to handle. At the extreme there is *linguistic phenomenalism*, which states that there is an equivalence between sentences about physical objects and some appropriate sentence about sense-data.¹⁴ The theory is put in terms of sentences because the positivists who propounded this theory did not wish phenomenalism to be an ontological theory. Positivists tend to be prejudiced against ontology, regarding very general questions about what sorts of things exist either as meaningless, or as questions to be settled by stipulation. Apparently ontological questions, therefore, are to be treated as requests for a decision concerning which vocabulary it is most convenient to employ.¹⁵ Preference was for the sense-datum vocabulary on the grounds that it is more inclusive than physical object vocabulary; it can be employed to report after-images, hallucinations, illusions and bodily sensations as well as veridical perceptions.

Linguistic phenomenalism has few supporters. First, the positivist attitude to ontology is not widely shared. Most philosophers believe, therefore, that if some translation of physical-object sentences into sense-datum sentences could be provided that would show that the physical world *in fact* consisted only of sense-data, not that two languages were equivalent. This could be expressed by saying that the different languages enshrine different theories and that theories contain ontologies: in preferring a language one is, therefore, choosing an ontology, and in choosing it one is committed to taking it seriously. Second, there is widespread scepticism about whether the proposed equivalences could ever be stated. No one ever seemed clear what the sense-datum equivalent of 'there is a table in the next room' could be. It would have to make no reference to public space or to any physical object. The superficiality of preliminary sketches such as 'if you go into the next room you will have table-type sense-data' only serve to emphasise the massiveness of the task in eliminating physical references. Third, even if such technical problems could be overcome, and it did turn out to be possible to say what types of experience go along with the truth of physical-object statements, it would not follow that they were equivalent. If one believes that our physical-object discourse is committed to the existence of a metaphysically mind-independent realm, then *no* set of facts about experience would be strictly equivalent to statements about physical objects; for any such statements could be true and everyone be

having harmonised hallucinations. This objection seems to presuppose that our ordinary physical object language is metaphysically realist. It is enough, however, that metaphysical realists *believe* that our language is realist for the availability of a material equivalence of physical object and sense-datum statements not to be enough to prove the truth of phenomenalism: for if *in fact* there is no harmonised hallucination and the statements about physical objects are, in fact, true if and only if the statements about sensedata are true, they would be materially equivalent but it would still not follow that there was no more to the existence of the physical world than what was reported in the sense-datum statements. The truth of the material equivalence 'something has a heart if and only if it has a kidney' does not show that having a heart expresses what it is to have a kidney.

Nevertheless, even a factual phenomenalist—one who believes their theory to be a matter of ontology, not language—is likely to feel drawn to some kind of analytical programme. They will accept that it is revisionary of our ordinary beliefs, in so far as they are metaphysical realist and that, therefore, the possibility of such an analysis is not sufficient to show the truth of phenomenalism. Rather, analysis is a necessary condition for the acceptability of phenomenalism, because without analysis we have no reason to believe that something that could be thought of as a physical world could be constructed simply out of actual and possible sense-data. There is behind this an extremely appealing principle which might be called 'the principle of reduction'. This is the principle that it ought to be possible to articulate how the basic entities that constitute a domain come to make true all the statements that are true of that domain. This, however, is not enough to require analytic reduction. It requires analytic reduction only if the only way that the basic ontology can intelligibly sustain any derived ontology is because one can intuitively see how statements couched in terms of those basic entities convey all the information contained in statements that refer to derivative objects. There may, however, be another way, according to which we can understand how, by some imaginative or analogical process, the basic ontology can be seen as, or interpreted as, constituting the derived ontology. I shall discuss what this other way might be on pp. 231–8. First, I shall bring a strong *prima facie* case for thinking that some form of phenomenalism must be intelligible.

5 The Berkeleian possibility

The simplest way of investigating the coherence of some form of idealism or phenomenalism is to consider the intelligibility of a straightforward kind of Berkeleianism. It requires the possibility of two things. The first is that there be subjects or minds the existence of which does not presuppose the existence of a mind-independent physical world. This boils down to the requirement that neither a materialist account of the mind nor a materialist account of personal identity is necessarily true, for if neither of these things is necessarily true then it is possible that minds be logically prior to the physical world. The second requirement is that it be possible that God should give to these minds experiences of the same kind as we have, and that they interpret them as being of a physical world, as we do, but there be no more that is real and existent about their physical world beyond the experiences God gives them and His intention to maintain the pattern in those experiences. If these two conditions are coherent, then at least one form of phenomenalism is coherent, even if it is not true of the actual world. Furthermore, once the Berkeleian possibility is shown to hold, then it looks as if what one could call the 'Humean possibility' is also available. For if a god could cause the appropriate experiences then it looks as if it should be at least a bare possibility that these experiences could occur spontaneously.

There are three points at which the 'Berkeleian possibility' might be attacked. First, someone might maintain that the very idea of a mind or subject is secondary to that of matter and material bodies. This is obviously not a question that can be properly discussed here, but my purported refutations of Wittgenstein's attack on privacy and of physicalist theories of perception are both relevant.¹⁶ Second, the concept of God might be thought problematic. For these purposes, supposed problems with the orthodox Judaeo-Christian conception of God are not relevant; any super-agent would do, and there are not likely to be problems with this once the possibility of immaterial minds has been conceded. The interesting objection to the Berkeleian possibility is the third. This is that a mind could not both (and perhaps not either) have experience like ours and interpret it physically if there were no more to the physical world than these experiences. The spirit of this objection is Kantian; our having experience presupposes an objective physical world, so that objective world cannot be constructed from experience.

This argument loses most of its force in the present context when a certain distinction is made. This is the distinction between (a) the claim that experience must be interpreted in physical-object terms and that physical-object concepts are primitive—that is, they are not definable or analysable in purely phenomenal terms: (b) the claim that we could not have the sort of experience we do unless there *actually existed* a mind-independent physical world. It is difficult to see how *philosophical* arguments—for example, arguments resting on conceptual analysis—could establish anything stronger than (a). A philosophical argument that proved (b) would be rather like the ontological argument for the existence of God: that is, it would be proving the existence of a non-abstract object by an *a priori* argument resting on no more than our concepts, and should be greeted with the same suspicion as greets the ontological argument. By contrast, because (a) essentially concerns the relations of concepts, it is not so improbable that there should be a philosophical argument for this conclusion. The broadly Kantian strategy can, at most, therefore, be used to establish the irreducibility of the physical interpretation—that is, the falsehood of analytical phenomenalism—and not to show that anything ultimately mind-independent answers to these concepts. This is, of course, why Kant called himself an *empirical realist* but a *transcendental idealist*: the first because the world we experience must be interpreted in physical terms, the second because there is no intelligible mind-transcendent physical reality.

D.M. Armstrong has an objection to phenomenalism which, if it worked, would refute the Berkeleyan possibility.¹⁷ Armstrong argues that the phenomenalist cannot accommodate physical time. He assumes that, for a phenomenalist, both physical space and physical time must be constituted out of the relations between phenomenal events. In the case of space, he thinks that it is at least barely possible that it could be constituted by the fit between different subjects' experiences, that is, on the basis of the way the qualitative and perspectival contents of our experiences can be harmonised. But he is resistant to the idea that our common time order is merely a function of the way our experiences fit together in their contents.

I think the phenomenalist could make either of two replies to this. First, they might deny that public time was a construct from phenomenal time: they might, that is, be prepared to allow that there is a common time in which all subjects live and which is basic and real, and that this is the time to which constructed physical objects are

attributed. Second, if the first option seems too feeble, they could ground physical time in the harmony of contents *plus* the causal relations that hold between the contents of different minds. Thus my experience of seeing my arm move occurs at the same time as your experience of seeing someone's arm (not your own) move, because there is a causal connection between my willing to move my arm and both my sense-data of an arm rising and your sense-data of a similar event. This answer presupposes that there is some sort of real causal connection between my sense-datum and yours: a Humean causal connection will not do, for a Humean conception of causation itself *presupposes* temporal ordering of cause and effect. It is difficult to see how there could be such a real connection between phenomena in different minds unless it were indirect. It is not plausible to suppose that my sense-data exercise a direct causal influence on yours, if only because it is difficult to see how there could be reasonably systematic causal laws of the right type. The laws that naturally come to mind would work via the intermediate physical world, but, *ex hypothesi*, the phenomenalist does not believe in the kind of physical world that could carry real, non-Humean causal powers. The indirect causal connection between our data would presumably have to be, therefore, via the activity of God. So the construction of time requires a Berkeleyian idealism. Of course, any theory that maintains that time is ultimately unreal outside experience is profoundly counter-intuitive, but this need not be an objection to phenomenism. If time is an ultimate constituent of reality then the phenomenist need not construct it, if it is not ultimate then phenomenism probably needs to be theistic. The only position that the phenomenist cannot tolerate is that time is ultimately real and that it presupposes physical reality, for then the physical world would have to be part of ultimate reality, which a phenomenist cannot maintain. But there is no reason why a phenomenist should feel obliged to maintain that time presupposes an unconstructed physical reality, so the problem should not worry them.

6 Humean, non-analytic reductionism

Most of the arguments brought against the coherence of phenomenism are arguments designed to show that physical concepts cannot be analysed in terms of or 'constructed from' or defined in terms of experiential contents, as a sense-datum theorist

would construe such contents. There are two interesting questions here. The first is whether such phenomenalist analyses can be carried out and the second is whether phenomenism has any reason to be committed to the possibility of such analyses. The Berkeleian possibility seems to show that, if one is interested in phenomenism as an ontological thesis then there need be no commitment to the possibility of reduction. The situation may be different if one is concerned to follow the empiricist objective of explaining the contents of all our non-formal concepts exhaustively, or as near exhaustively as possible, in terms of the contents of experience. A phenomenalist analysis would certainly meet this requirement, but is it the only account that would do so? Despite what appears to be his strict empiricist theory of meaning, Hume's account of how we come to think of the physical world is not analytically reductive.¹⁸ It is not a matter of defining physical concepts in phenomenal terms, but of imagining the physical world using only the materials provided by experience, phenomenistically conceived. Hume's idea is that we move from private impressions to the physical world, not by an unconscious analytical construction but by a spontaneous imaginative leap. The very regularity of our visual impressions leads us to treat them as if they had a life of their own and existed irrespective of their being perceived. We, therefore, think of them as mindindependent objects available to all perceivers.

One might characterise Hume's psychological account as genetic rather than analytic and then worry that a purely genetic account will not do the job required. The account is genetic because it explains how we come to think of impressions as something they are not. It might seem that a genetic account is not enough because we are trying to explain how content of one kind—thought about physical objects—is extracted from content of another kind—apprehending sense-data—and in this context relations between cognitive contents must be a matter of the relations of something rather like meanings and how one kind of content constitutes or explains another is not a matter of brute causal fact but of something more like an analytical relation.

It is true that a purely genetic account could not be relevant, but a purely definitional account does not allow for the role of analogy and imagination. The imagination operates spontaneously according to Hume's genetic account and analogy is appealed to in the rational explanation of how the imagination works. Because experience does

not present its objects to us as being dependent for their existence on our experience of them, we can imagine that they may exist when we are not perceiving them: the world when unperceived can be conceived of as analogous to the world as we perceive it. If vision were like pain, where the object of experience is presented as essentially dependent on its being experienced, then this analogy would not be possible. Similarly, we can conceive of a spatial realm that is not confined to our sense-fields because there is nothing in our visual experience that blocks the idea that its visual dimensions are extended beyond it. Indeed, the very natures of space and time as they are presented seem to force upon us the intelligibility of their recursive extension beyond the contents of the experiences in which they are given. It seems more difficult to get one's imagination round the idea that space ends at the edges of one's visual field, or that time is confined to one's experiential contents, than it is to project them indefinitely in all directions.

It is true that Hume would not see space in this last way. Hume would not allow that there is anything in the nature of space as it is given in an individual perception that suggests its extendability, for it is against his principles to allow that an impression can carry any implication about anything outside itself; all such things are learnt by association. For Hume, therefore, it is the different kinds of ordering that visual experiences have, contrasted with pains and sensations—namely that the latter ‘move around with us’ and the former do not—that creates the illusion of externality and not something about the individual experience. It would, therefore, make no sense, for Hume, to suggest that there could be experiences which did not seem to us to be of external things, but which it would be just as useful to explain by reference to external causes as it is with our visual impressions. The very fact that this explanation works without undue complexity is what makes the objects seem external. His position is still different from an analytical one, for it is not that certain types of order analytically constitute externality, but that they make their contents *seem* external: the difference is phenomenological, not constructed.

These disputes about Hume are not essential to the general point. *Some* account of the nature of experience is required which accounts for our possession of the idea of externality. The fundamental difficulty with the theory is the same in its more or less nominalist forms, namely that it seems to be an error theory. By equating statements about

physical objects with statements about sense-data, the latter being true, analytic phenomenalism preserves our physical object beliefs as true. The Humean kind of construction, on the other hand, works only because we mistakenly read our experiences in a naive realist manner. A representative realist might accept Hume's account as true of the development of our concepts and not be worried by this, for this would not commit them to saying that our conception of the world was essentially mistaken. For the representationalist, the analogies and imaginative projections that Hume describes are what enable us to be right about the physical world, though we are wrong in thinking that our sense-data are the world itself; for them, the physical world actually possesses certain crucial features that we wrongly read into the phenomenal world. For the phenomenalist, however, nothing actually possesses those features: if we did not make mistakes about the real nature of experience, we would never arrive at the conception of a physical world, which is a conception to which nothing strictly answers.

This is not, however, the correct way of reading the neo-Humean construction. It is not by some mistake that we come to see the world as external, rather it is inevitable and unavoidable. Our naive picture of the world is a way of construing experience and, as such, is neither true nor false, but successful or unsuccessful. It is manifestly successful for all ordinary purposes of living. This is not to say that our naive conception is a matter of choice or convention: on the contrary, it is the only conceivable (probably) way in which we could make sense of the flow of our experience. It is, therefore, the correct way of reading our experience and the way that experience is meant to be read. It becomes false only if we become naive *realists* and think of the naive world as something more than the world of experience.

It is natural to regard this approach as disingenuous. Surely realism is built into our naive understanding of the world and is not merely a false philosophical gloss on it. There is some truth in this, but not enough to be damaging. The thought that the idealist reading of our naive grasp on the world is insufficiently realist rests, in part, on mistaken attempts to picture it. It is pictured as representing the world as a harmonised hallucination shared by millions of brains in vats: within a real space each subject lives within a false imaginary space. The idealist theory is, therefore, assimilated to dreams and hallucinations as we actually know them. But suppose that there is no

space (and perhaps no time) outside the experience of the minds, then it becomes impossible to imagine the theory as a kind of collective solipsism. Indeed, it becomes impossible to *imagine* the theory at all: it makes no sense *to picture* idealism as a metaphysical theory: it purports to be a philosophical truth and to judge it finally one must conceive of it only intellectually.

Imagination conspires against idealism as an account of common sense, but this conspiracy is much weakened by the remembrance that naive realism is no longer an option. Idealism, therefore, has only to establish its credentials as a better articulation of our normal beliefs than representative realism. (It is in this context that Berkeley's claim to be defending common sense should be seen by a charitable reader. He is presenting a perspective for rescuing common sense, given what science and philosophy have shown.) Two points can be made in this connection. The first is that the 'brain in vat' image of the individual's experience is much more applicable to representative realism than to idealism, because, according to the representationalist, the individual really is located in a public space, enclosed in a bubble of private experience. The second is that, as Hume pointed out, the claim of realism to be common sense entirely rests on naive realism and cannot legitimately be transferred to representative realism:

[representative realism] has no primary recommendation either to reason or the imagination, but acquires all its influence on the imagination from the former [i.e. naive realism].¹⁹

In other words, the seemingly overwhelming intuitive priority of realism entirely draws its force from our feeling that the immediate objects of our awareness are mind-independent physical objects. This is the sole origin of the psychological power of realism. Once one is persuaded that these immediate objects are not mind-independent, then the intuitive priority we give to realism has entirely lost its rationale. It is bad faith to pretend that representative realism answers to the intuitions that make realism psychologically compelling, because to get to the representative level we have already discounted as false the substance of that intuition, which was that the immediate objects of awareness are mind-independent. This is surely one matter on which Hume's sceptical argument rings true. Once naive realism is refuted, the explanation of the nature of the physical world is entirely open and no one type of account starts with any greater claims on our credence than any other.

Against this it might be argued that a theory which makes out that the origin of our experience is rather like the world which naive realism posits, is inherently more plausible than a theory which makes the origin of our experience very different from what we had presupposed. Now such a theory is likely to go some way towards satisfying our psychological urge towards realism, but that does not mean that it answers the intuitive ground or rationale of that urge. To see this, suppose that all our experience seemed to have immediate objects, so that we were no more naturally inclined to think of our visual, tactile, auditory sense-data as external than we are to think so of our pains. Suppose, further, that we were nevertheless still prone to look for an explanation of our experience and of why it fitted into such useful patterns. In these circumstances a representative physical realist explanation would not commend itself intuitively any more than the alternatives. Indeed, we would probably be inclined to doubt whether the notion of something rather like these phenomena yet outside the mind really made sense, just as we think it obvious that nothing outside the mind could resemble pains or other bodily sensations. It is surely not plausible to maintain that representative realism is made genuinely more credible by its fitting a *mistaken* understanding of our experience (namely our inclination to take the immediate objects of our experience to be independent of the mind) when, if we were not prone to make this mistake, it would seem especially implausible amongst the accounts of the origin of experience.

7 Phenomenalism and common sense

The burden of these arguments is that common sense can be divorced from realism. The things that we naively say can be deemed correct if they are read as part of our constructive interpretation of the world, rather than as reporting how things ultimately are. The naive picture of the world is constrained by certain canons of interpretation that bear some of the marks of being conventions. For example, under the circumstances described by the argument from illusion, when an object looks other than it really is, we say that it seems F though it is G and put the burden of the distortion into our act of seeing, pretending that the ostensive object of our awareness remains the same. This is because, construed as a mind-independent object, it has remained the same, and, within that framework, the distortion is to be attributed to the seeing. This is part of the 'logic' or 'grammar' of our naive physical

interpretation of experience. But, for reasons given earlier in this book, this is a matter of how we must talk and think about experience and is not accurate as an analysis of the phenomena. This does not make the naive interpretation false, for that interpretation is not an analysis of the phenomena, but an interpretation of it: it is not saying what is ultimately there or giving the correct metaphysical perspective, but articulating the phenomena in the way required for it to constitute a medium for living. The logic or conceptual structure of naive realism requires phenomenalism (and, hence, the abandonment of the 'realism' component as usually understood) because if it is taken in the usual realist way it has no account within its ontology for non-veridical phenomena ('illusions') and no plausible understanding of the relation between experience and the causal processes of perceiving. These facts have been shown by the failure of attempts to save naive realism by appeal to intentional objects and by the success of the modified causal argument for sense-data. Only the phenomenalist can give a coherent account of intentionality for sensations. A content is intentional if, despite being phenomenally real, it is not deemed to be realised in the world as naively conceived. This is ontologically unproblematic; being left off the first team does not endow someone with a mysterious ontological status.

In the light of this we can look again at the disjunctive theory of perception. I argued that the disjunctive theory was totally implausible for it attributed variable effects to the same proximate cause.²⁰ The principle 'same proximate cause, same immediate effect' is essentially a realist principle: if it were a matter of *stipulating* things to be causes and effects there would be no reason to adhere to it. But we now see the naive picture of the world as being a necessary way of construing experience, whilst not being true in a realist sense. Just as we need to deal with the phenomena of illusion within a grammar that is naively realist, so we can do the same for hallucination and perception. We simply deem the perceptual cases to be different from the hallucinatory because that is required to make the naive reading of perceptual experience: this, too, is part of the 'grammar' of our naive understanding of the world. The phenomenalist is not absolutely obliged to adopt the disjunctive theory. They could, within their constructed physical world, allow the causal argument for sense-data to go through. If they did so they would not face the same problems as the *realist* representationalist. It would not be a mystery for them what matter really consisted of, for, *ex hypothesi*, they would adopt the

naïvest form of representationalism, with the physical world possessing all sensible qualities. Nor would there be a problem about whether the world beyond the veil instantiated the laws of physical science, for that it did so could be stipulated in its construction. It is difficult to see, however, what motive a phenomenalist would have for taking this view. If they were a practising neurologist they might find it natural to adopt the principle of same proximate cause, same immediate effect when in their laboratory. But, as there would be no fact of the matter, they should follow the disjunctive theory for other purposes. Not being a realist about the physical world, the relations between the brain and experience do not represent simple causal facts but the interface of two domains, one of which is strictly real and the other a necessary frame of interpretation. There is, therefore, scope for decision about how to characterise their relation and the easiest decision would be the disjunctive one.

8 Conclusion

We have seen that representative realism resists the usual objections made by the naïve or direct realist, but looks more vulnerable to some of the criticisms of the phenomenalist. Phenomenalism, it seems, is coherent, provided it does not take the analytical form and, indeed, is the only way of rationalising the discourse of naïve realism, naïve realism itself having been decisively refuted in earlier chapters.

Notes

I The Classical Empiricist Conception of the Content of Perceptual Experience

- 1 The term 'sense-datum' was introduced earlier this century with a slightly different sense from that which I have given it. According to Moore or Price (as we shall see in Chapter II) a sense-datum is whatever qualitative object we are aware of in perception. Thus, in Price's famous case, when we see a tomato we are aware of a red bulgy patch: that is a sense-datum. The question can then be raised as to whether this datum is or is not part of the surface of an external, physical object. The answer to this has to be 'no', mainly for reasons given in the 'argument from illusion'. Once one accepts that whenever one seems to see, for example, a red patch, then there really is, somewhere, a red patch of which one is aware (what I call in Chapter II 'the Phenomenal Principle') then the success of the argument from illusion is a foregone conclusion. This principle is explicit in Moore and Price's notion of sense-datum; it is implicit, that is, that there really is a red patch whenever there seems to be, for 'sense-datum' is simply the name for such objects. So the definitionally neutral sense-data turn out in fact to be sense-data in my sense (or more or less). This is why most people think of sense-data in the way I have defined the notion.
- 2 For Russell, see *Analysis of Matter*, and Lockwood's commentary in *Mind, Brain and Quantum*, 156ff. For the phenomenologists, see pp. 21–7.
- 3 Aristotle's account of perception is mainly found in chs. 6–12 of book II of *De Anima* and in *De Sensu*. It is not universally agreed that phantasms are involved in ordinary perception for Aristotle—they are associated with imagination in the ordinary sense of that term. Nevertheless, Aristotle says that imagination consists in a weakened version of the disturbance that causes perception (*De Anima*, 428b26f), so if there is a physical image in imagination, there must be a stronger one involved in perception. The scholastics uncontroversially have them involved in perception. For a useful introduction to the problems in ancient and medieval philosophy, see the first three chapters of Hamlyn, *Sensation and Perception*.

- 4 Empedocles and Democritus seem to have had theories of this kind. The main source for these and other early theories of perception is Theophrastus's *De Sensu*.
- 5 Jonathan Barnes has pointed out in discussion that Aristotle never explains how an image can carry a substantial form. Like a sense-datum, it can possess sensible qualities—the shape and colour of Socrates—but not the form of his humanity. One could try arguing that it bears this intentionally, but this does not fit with the idea that a form becomes intentional by being *abstracted* from the image, where it is literally present. The implication is that Aristotle is too careless about epistemological questions and does not see the problem that his account of the machinery of perception poses for the acquisition of substantial forms.
- 6 The view that Aristotle's materialism was traduced by his non-materialist commentators is to be found in Sorabji, 'From Aristotle to Brentano: the development of the concept of intentionality' and 'Intentionality and physiological processes: Aristotle's theory of sense perception'.
- 7 Brentano, *Psychology from an Empirical Standpoint*, 88.
- 8 Cohen, 'St Thomas Aquinas on the immaterial perception of sensible forms'. There is a reply in Haldane, 'Aquinas on sense perception'.
- 9 *De Anima*, III, 4.
- 10 The case for individualised forms is found in A.C.Lloyd's *Form and Universal in Aristotle*: the case against is stated by Michael Woods in 'Universals and particulars in *Metaphysics*'.
- 11 See, for example, the proof of the existence of God in the third *Meditation*.
- 12 Locke, *An Essay Concerning Human Understanding*, I, i, 8.
- 13 *Ibid.*
- 14 In his attack on abstract ideas in the introduction to *The Principles of Human Knowledge*.
- 15 *A Treatise of Human Knowledge*, bk. 1, pt. 1, sec. 1.
- 16 The discussion of externalism in Chapter V, pp. 136–50, is relevant here.
- 17 This seems true, for example, of Yolton's response to the fact that Locke calls ideas 'perceptions'. See note 18 below.
- 18 The exception being John Yolton in *Perceptual Acquaintance*, a book dedicated to showing that Locke, Berkeley and Hume all held to an intentional conception of ideas, and that we think otherwise only because Reid foisted a misinterpretation of them on us. That this is almost certainly wrong for Locke is demonstrated by Ayers in the first volume of his majestic *Locke*, 44–69: if it is not true of Locke it is morally certain that it is not true for Berkeley or Hume, for whom it never was plausible.
- 19 For Descartes' view of matter see, for example, *Principles of Philosophy*, pt. II, sec. 4; for a thorough discussion of how this affects his account of the relation between appearance and reality, see Chapter I of Hacker's *Appearance and Reality*, especially 8ff. For Locke's account of the difference between primary and secondary qualities, see *Essay*, II, 8. It is the principal thesis of Hacker's book that a false division between primary

- and secondary qualities is what produces the (he believes) false division between appearance and reality and the sense-datum theory of perception.
- 20 This difficulty may be somewhat lessened if we follow Cohen (see note 8 above) in assimilating the physicality of the phantasm to that of reflections and mirror images. Perhaps if they are allowed to be both intentional and physical they can possess secondary qualities. As the problems with the mind-independent status of mirror images seem insuperable, to attenuate their physical character further by adding qualities that matter does not possess, seems to make the situation worse rather than better.
 - 21 This could be got from Descartes, *Rules for the Direction of the Mind*, Rule 12, 40–1 in vol. 1 of *Philosophical Writings*: and from Locke's remark that secondary qualities are powers of primary qualities 'when they operate without being distinctly discerned' (II, viii, 22). I can see no reason to think that either Descartes or Locke intended to impute the vagueness to the *experienced* idea, only to what it failed to reveal about the underlying primary structure. Descartes, indeed, explicitly combines these thoughts: 'when we say that we perceive colours in objects, this is really just the same as saying that we perceive something in those objects whose nature we do not know, but which produces in us a certain very clear and vivid sensation which we call the sensation of colour' (*Principles*, I, 70, in *Philosophical Writings*: quoted in Ayers, vol. 1, 39).
 - 22 First *Meditation*.
 - 23 This argument will be properly developed in Chapter VI. Although this argument is implicit in most Cartesian and empiricist thinking, I am not aware that it was produced as an argument before C.D.Broad did so in 1952.
 - 24 A.D.Smith, 'Berkeley's central argument against material substance'.
 - 25 Hacker's verdict on Reid, in *Appearance and Reality*, 42–3, is essentially correct.
 - 26 The quotations from Reid are from *Inquiry into the Human Mind* in *The Works of Thomas Reid*. This reference, 137.
 - 27 *Ibid.*, 138.
 - 28 *Ibid.*
 - 29 *Ibid.*, 140.
 - 30 *Ibid.*, 137–8.
 - 31 *Ibid.*, 136.
 - 32 *Ibid.*
 - 33 Brentano, *Psychology from an Empirical Standpoint*, 92.
 - 34 *Ibid.*, 93–4.
 - 35 W.J.Ginnane, 'Thoughts'.
 - 36 Brentano, *Sensory and Noetic Consciousness*, 5, my italics.
 - 37 Brentano's failure to make this distinction becomes more intelligible if we accept David Bell's interpretation of him (Bell's *Husserl*, ch. 1). According to Bell, Brentano treats all mental contents as presentations, thus making problems for a round square in thought as well as in perception. What one

affirms in a judgement is a presentation of a state of affairs, not a proposition. This is just the opposite of what one might expect from someone who emphasises intentionality and takes the Aristotelian account of the intentional reception of form as his model. It is worth noting, too, that by 'inexistence' Brentano meant only existence *in* the mind—it does not signify a weak kind of existence which, somehow, accommodates the *now*-existence of some intentional objects. There are certain features of Bell's interpretation of Brentano that are difficult to square with some of the passages I quote, but both his and my evidence support the view that Brentano had no clear idea of how to explain intentionality, nor of how to apply it to perception in a way that would assist perceptual realism. My discussion of Brentano and Husserl fits more easily into the problem as construed by Føllesdal in 'Brentano and Husserl on intentional objects and perception'.

- 38 See Føllesdal's 'Husserl's notion of *Noema*', 74. This article gives a very lucid account of the idea of *noema*. In Husserl himself, discussion of *noemata* is to be found throughout the *Logical Investigations* and *Ideas*.
- 39 This 'X' is to be found, for example, in *Ideas*, sec. 131. See D.W.Smith, 'Husserl on demonstrative reference and perception'.
- 40 See Føllesdal, 'Brentano and Husserl on intentional objects and perception', 39.
- 41 Roderick Firth, 'Sense-data and the percept theory', vol. 1, 447–8.
- 42 *Ibid.*, 449.
- 43 H.H.Price, *Perception*, 63.

II The Traditional Arguments for the Empiricist Conception of Sense-contents: the Argument from Illusion

- 1 J.L.Austin, *Sense and Sensibilia*, 49.
- 2 For classical scepticism, see Annas and Barnes, *The Modes of Scepticism*, and Barnes, *The Toils of Scepticism*.
- 3 Hume, *Enquiries*, 151.
- 4 *Ibid.*, 152.
- 5 *Ibid.*, 152.
- 6 *Ibid.*, 152–3.
- 7 *Essay*, II, xiii, 21.
- 8 Broad, *Scientific Thought*, 238, reprinted in Swartz, *Perceiving, Sensing and Knowing*, 88. The italics are mine. Swartz is an excellent collection for twentieth-century classics in the philosophy of perception, and I shall give references for items collected there.
- 9 *Ibid.*, 239 and 89.
- 10 Moore, 'Some judgements of perception', 21–3; Swartz, *Perceiving*, 20–2.
- 11 *Ibid.*, 23–4; Swartz, *Perceiving*, 22–3.
- 12 *Ibid.*, 24 and Swartz, *Perceiving*, 23.
- 13 Moore, 'Visual Sense-data', 208; Swartz, *Perceiving*, 133–4.

- 14 Price, *Perception*, 3.
- 15 J.M.E. McTaggart, *The Nature of Existence*, ch. 23; Bradley, *Appearance and Reality*, esp. ch. 3.
- 16 W.F.H. Barnes, 'The myth of sense-data', reprinted in Swartz, *Perceiving*, 138–67.
- 17 Ibid., 104 and Swartz, *Perceiving*, 153.
- 18 Ibid., 106; Swartz, *Perceiving*, 155.
- 19 Price, *Perception*, 55.
- 20 Ibid., 56.
- 21 Ibid., 57.
- 22 Ibid., 57.
- 23 Price, *Perception*, 55–8; Ayer, *Foundations*, 14–21. I say that Ayer defends this option 'in a sense' because he admits that, though a kind of direct realism can be preserved in this way, it involves revision of our naive concept of body. This would appear to leave the naive realist with a choice between abandoning our naive conception of the directness of perception or abandoning our naive conception of its object. Ayer tries to avoid this by making two moves which he conflates. First, he says that the naive realist can appeal to intentional idiom, and, second, that this differs only linguistically from the sense-datum language. In fact, if the intentional idiom were adequate as an elaboration of naive realism, there would be no problem and the sense-datum language would be redundant. I discuss Ayer's later, but essentially similar, views in Chapter IX.
- 24 *Perception*, 35.
- 25 Ibid., 62.
- 26 Ibid., 63.
- 27 Ibid., 63; my italics.
- 28 Ibid., 64.
- 29 I have chosen a primary quality, not a secondary, to illustrate the point, because it is an open question whether reference to appearance enters into secondary qualities, whereas it clearly does not for primaries.
- 30 The reduction of perception to judgement is discussed in Chapter V.
- 31 Broad, *Scientific Thought*, 236; and in Swartz, *Perceiving*, 87.
- 32 Ibid., 236–7 and Swartz, *Perceiving*, 87–8.
- 33 Both quoted remarks are from Austin, *Sense and Sensibilia*, 26.
- 34 Grice, 'The causal theory of perception', reprinted in Swartz, *Perceiving*.
- 35 Hacker, *Appearance and Reality*, 33–4, my italics.
- 36 Ibid., 223.
- 37 Ibid., 223.
- 38 Broad, 'Some elementary reflexions on sense-perception', 9, and in Swartz, *Perceiving*, 36–7.

III Further Arguments against Naive Realism

- 1 Kirk, Raven and Schofield, fragment 589.

- 2 From *Il Saggiatore*, excerpted in Danto and Morgenbesser, *Philosophy of Science*, 27–8.
- 3 Descartes, *Philosophical Writings*, ‘Replies to the sixth set of objections’, II, 297.
- 4 Locke, *Essay* II, viii, 15.
- 5 Ibid. II, viii, 17.
- 6 Hacker, *Appearance and Reality*, 139–40.
- 7 Armstrong, *A Materialist Theory of the Mind*, 270–90. The parallel with the topic-neutral account of mind is obvious. In both cases there is a need to eliminate qualitative content to facilitate a reductive physicalist account.
- 8 For an account of the resemblance theory of universals, see D.M. Armstrong’s *Universals and Scientific Realism*, vol. 1, 44–57.
- 9 Hardin, ‘Colour and illusion’, in Lycan, *Mind and Cognition*, 555–67.
- 10 Price discusses the ‘selective theory’ in *Perception*, 40–53, and Hirst has a useful section on selection versus generation in *The Problems of Perception*, 62–6.
- 11 This is how I understand John McDowell’s theory in, for example, ‘Values and secondary qualities’.
- 12 See Chapter IV for privacy and pp. 167–72 for sensation and grammar.
- 13 Locke, *Essay*, II, xxiii, 11.
- 14 Eddington, *The Nature of the Physical World*, xi–xiii; quoted in Stebbing, *Philosophy and the Physicists*, 47–8.
- 15 This argument is used, for example, by Stebbing, *Philosophy*, 45, when discussing Eddington.
- 16 Dummett, ‘Common sense and physics’; Strawson, ‘Perception and its objects’.
- 17 Strawson, ‘Perception’, 57–8.
- 18 Ibid., 58.
- 19 Descartes, *Philosophical Writings*, ‘Principles of Philosophy’, I, 284. Quoted by Hacker, *Appearance and Reality*, 15.
- 20 Locke, *Essay*, II, viii, 19.
- 21 *First Meditation*.

IV Sense-data and the Anti-private Language Argument

- 1 *Philosophical Investigations*, I, 246 is probably the nearest one can get to the explicit statement that one cannot know that one is in pain: it follows from this and the immediately succeeding paragraphs that there is nothing to know.
- 2 *Philosophical Investigations*, I, 293.
- 3 See Chapter V for an account of such theories.
- 4 Hopkins, ‘Wittgenstein and physicalism’, esp. 122ff.
- 5 *Philosophical Investigations*, I, 265.
- 6 A.J. Ayer, ‘Can there be a private language?’, 41–3.
- 7 Hacker, *Appearance and Reality*, 225.

- 8 For a convincing revision of the notion of incorrigibility, see Richard Warner, 'Incorrigibility'.
- 9 Kripke's argument is the content of *Wittgenstein on Rules and Private Language*.
- 10 *Philosophical Investigations*, I, 139ff.
- 11 *The Concept of Mind*, 113ff.
- 12 Kripke, *Wittgenstein and Rule-Following*, 8–9.
- 13 *Ibid.*, 23ff.
- 14 Of course, it does not follow from the intellectualist conception that we do not make mistakes. How we can make mistakes even though we have a direct intellectual grasp on what we intend is something that would be explained in an adequate account of embodiment. I attempt a beginning of this in 'A dualist account of embodiment'.
- 15 Kripke, *Wittgenstein and Rule-Following*, 51.
- 16 *Ibid.*, 51.
- 17 *Ibid.*, 54.
- 18 Goodman, *Fact, Fiction and Forecast*, ch. 3, sec. 4.
- 19 Walker, *The Coherence Theory of Truth*, 142–5.
- 20 *Coherence*, 142–3.
- 21 There is a general tendency for those who attack the 'given' in its traditional form, on the grounds that everything is theory-relative, to introduce surreptitiously their own surrogate. In behaviouristic theories, for example, the nexus of stimulus and response is treated as a completely objective given. Walker shows why this happens: without an implicit given, no account can get started.

V Contemporary Physicalist Theories of Perception

- 1 Place, 'Is consciousness a brain process?'; this reference, Borst, *The Mind-Brain Identity Theory*, 48–51.
- 2 This is the theory of Smart's classic paper 'Sensations and brain processes', and of Armstrong's *A Materialist Theory of the Mind*.
- 3 For arguments that it is an illusion that intellectual states can plausibly be handled in a physicalist way, see the first six chapters of *Objections to Physicalism*, ed. Robinson.
- 4 *A Materialist Theory of the Mind*, 208.
- 5 The argument of this section appears in my *Matter and Sense*, ch. 6, sees. 2 and 3.
- 6 Feigl, 'The "mental" and the "physical"', 433.
- 7 Churchland, *A Neurocomputational Perspective*, 23.
- 8 *Ibid.*, 30.
- 9 *Ibid.*, 30.
- 10 *Ibid.*, 31.
- 11 Dretske, *Knowledge and the Flow of Information*, 143.

- 12 I do not want to suggest that the protagonists of analogue content fail to realise that they are functionalists: nevertheless, there seems to be an implication—or, at any rate, it is easy to mistakenly draw the implication—that something is being added by using the notion of analogue content which makes the account less phenomenologically counter-intuitive than neat functionalism. On this level, nothing is being added at all.
- 13 I have twice before tried to state the argument rigorously; once in ‘The flight from mind’ and, more recently, in ‘The anti-materialist strategy and the knowledge argument’. The development in the argument was not caused by discovering errors in the previous versions, but in the hope of making it more perspicuous.
- 14 Burge, ‘Individualism and psychology’, 32; quoted by Davies, ‘Individualism and perceptual content’, 464. Michael Martin has pointed out to me that Burge holds that the causal relation is a necessary and not a sufficient condition for being content; he does not, therefore, subscribe to a completely causal *analysis* of content. At proof, I noticed that my strategy against externalism is anticipated by Sheehan, ‘Aquinas on intentionality’. This excellent article is also relevant to medieval theories of perception.
- 15 P.Smith, ‘On “The objects of perceptual experience”’, 194–5.
- 16 Martin made these objections at a meeting of the Ockham Society in Oxford.
- 17 I believe that this is an accurate reconstruction of objections made by John Kenyon, Hugh Rice and Ralph Walker.
- 18 Armstrong does this in *A Materialist Theory of the Mind*, ch. 15:1 reply at greater length in *Matter and Sense*, 53–8.
- 19 C.Macdonald, ‘Weak externalism and mind-body identity’.
- 20 Kripke, *Naming and Necessity*, 81–2.

VI The Revised—and Successful—Causal Argument for Sense-data

- 1 This argument is used by Broad to refute naive realism in ‘Some elementary reflexions on sense-perception’. He does not take it as proving a full sense-datum theory, however, because he prefers an adverbial account of content. More of this in Chapter VII.

I claim that this argument is not clearly in the empiricists, nor, so far as I know, is it common currency before Broad. It has recently come to my notice that it is in the first of Malebranche’s *Dialogues on Metaphysics*:

Now on the supposition that the world was annihilated and that God nevertheless produced the same traces in our brains, or, rather, presented the same ideas to our minds which are produced in the presence of objects, we should see the same beauties...

But what I see when I look at your room...will still be visible even should your room have been destroyed and even, I may add, if it had

never been built! I maintain that a Chinese who had never been in the room can, in his own country, see everything I see when I look at your room provided—which is by no means impossible—his brain is moved in the same way mine is when I now consider it.

Malebranche's conclusion, however, is not a normal version of the sense-datum theory, for he holds that 'The dimensions I see are immutable, eternal, necessary'. His eccentric platonising of ideas, even as they constitute the contents of perception, and the absolute (though inadequately explained) distinction he makes between ideas and sensations, makes it difficult to categorise his theory.

- 2 A classical source for the sufficiency of brain stimulation for experience is Penfield's *The Excitable Cortex in Conscious Man*.
- 3 Don Locke, *Perception and our Knowledge of the External World*, 111–12.
- 4 *Principles of Human Knowledge*, sec. 8.
- 5 *Seeing and Knowing*, 71.
- 6 J.M.Hinton, *Experiences*, 75–93; G.Pitcher, *A Theory of Perception*, 54–7. The disjunctive theory is also discussed by Paul Snowdon in 'Perception, vision and causation' and by John McDowell in 'Criteria, defeasibility and knowledge'. McDowell is advocating the theory, but shows no sign of recognising that the functioning of the same proximate cause might constitute a problem for the theory. Snowdon, however, claims only to be analysing our ordinary concept and is, therefore, committed to no particular view about what to say on the basis of scientific hypotheses such as ours concerning the causal powers of the brain.
- 7 *Theory of Perception*, 56–7.
- 8 *Experiences*, 75.
- 9 *Ibid.*
- 10 *Ibid.*
- 11 *Ibid.*, 77–82.
- 12 See, for example, Davidson's approach to the role of interpretation in the construction of our cognitive states, and Dennett's instrumentalist construal of intentional systems in his early and middle writings.
- 13 When discussing phenomenalism in Chapter IX we shall see that the disjunctive theory can be given a role in that kind of non-realistic theory, for we can use it to guide our account of the relation of the mental to the physical in that context. This is because, within a phenomenalist constructed physical world, there is no fact of how the mental and the physical are related, it is a matter of how it is convenient to construct that relation. This is quite different from making the nature of the mental a matter for choice or convention within a physical realist framework.
- 14 Once again, the phenomenalist can be more flexible here, as we shall see in Chapter IX, p. 136–8.
- 15 This view is defended by Hinton in 'Visual experiences'.

VII The Intentional and Adverbial Theories

- 1 I noted in Chapter I that it has been argued that Brentano took *in* existence as signifying existence *in* the mind, rather than as signifying a special kind of quasi-existence. These two approaches are not so different if one adds to the notion of mental existence the proviso that the mental existence of F is not an instance of F: in this way the ontological uncertainty is imported into Brentano's analysis, for it is less than obvious in what way F characterises its mental existence.
- 2 I am not saying that Henry *really is* intentionally present in the picture in the way that an object is really present in thought. A scholastic who thought that mirror images were real cases of intentional existence (see Chapter I, pp. 7–9) might say the same about pictures. Others might make pictures a matter of interpretation, which cannot, on pain of regress, be the case for mind. (See *Objections to Physicalism*, ed. Robinson, 6–8.)
- 3 Anscombe, 'The intentionality of sensation', 164.
- 4 Ibid., 162.
- 5 Ibid., 164.
- 6 Ibid., 169.
- 7 This is implicit in the text, especially 166.
- 8 Ibid., 166.
- 9 Ibid., 162.
- 10 Ibid., 162ff.
- 11 Ibid., 162.
- 12 Ibid., 164.
- 13 Come to think about it, this is exactly Dennett's theory of consciousness in *Consciousness Explained*; consciousness does not exist—it is just subjectively as if it did.
- 14 Chisholm, *Perceiving: A Philosophical Study*, 120–5.
- 15 This is not an objection to materialism. The point is not that mental acts could not be physical events, but that a naive realist could not claim that the physical objects in the manifest image of the world are modes of mental activity, as opposed to objects of it.
- 16 William Lycan, *Consciousness*, 88.
- 17 Mainly in 'On the adverbial analysis of visual experience'.
- 18 In 'The adverbial approach to visual experience'. He introduces the adverbial operator 'Coin', so that 'S senses redly Coin squarely' expresses the relevant coincidentality of the sensings. He also invents an operator to express the separateness of contents. 'S senses redly Sep. roundly' is meant to express what is required for the distinctness of objects.
- 19 Ducasse, 'Moore's refutation of idealism'.
- 20 'Moore's refutation', 228ff.
- 21 In his replies in Schilpp, *The Philosophy of G.E. Moore*, 653–60, a volume dedicated to him.
- 22 See, however, the next chapter for another—but no more successful—interpretation of intentionality and perception.

VIII The Nature of Sense-data

- 1 The distinction between what is interpretative and what is intentionally but really present will emerge in the discussion of depth, pp. 205–7.
- 2 Armstrong, *A Materialist Theory of the Mind*, 218.
- 3 Ibid., 219.
- 4 Chapter V, p. 122f.
- 5 Armstrong, *A Materialist Theory*, 219–20.
- 6 Ibid., 220.
- 7 Dennett, *Consciousness Explained*, 117–18.
- 8 Ibid., 118.
- 9 Ibid., 118–19.
- 10 Ibid., 119.
- 11 Ibid., 114, quoting Goodman, *Ways of Worldmaking*, 73.
- 12 For a discussion of the inadequacies of Dennett's 'interpretational' account of the mental, see *Objections to Physicalism*, 6–8.
- 13 I attempt to discuss questions in this area in 'A dualist account of embodiment'.
- 14 Frank Jackson, 'Epiphenomenal qualia'.
- 15 Dennett, *Consciousness Explained*, 399–400.
- 16 Paul Churchland claims that we could be taught to respond to stimuli by scientific descriptions and that we would thus be 'directly aware' of micro-physical processes. (See, for example, *A Neurocomputational Perspective*, 54ff.) I have argued that there are general problems with this ('The anti-materialist strategy and the knowledge argument', 169–70). But even if there are no general problems, this idea cannot be used to justify Dennett's attribution to Mary of immediate recognition of physical inputs, for such teaching is available only to those who have the appropriate experiences. Mary has not previously experienced colour, so she could not yet have learned how to respond spontaneously to colour experiences with scientific descriptions.
- 17 See 'Functionalism and qualia' and 'Absent qualia are impossible: a reply to Ned Block'.
- 18 Jackson's belief that his *qualia* are epiphenomenal is based on his belief that the physical system is closed and can accommodate everything other than *qualia*. Dennett's criticism of Jackson's epiphenomenalism is, I think, powerful. Ayer combines *qualia* with behaviourism for other states in *Origins of Pragmatism*, 173–9. I discuss Ayer's position in *Matter and Sense*, 105–7.

IX Sense-data and the Physical World

- 1 Unless indicated, directly or by context, I shall use the terms 'phenomenalism' and 'idealism' interchangeably.
- 2 Chapter III, pp. 59–74.

- 3 *Principles*, sec. 10.
- 4 See Chapter VIII, pp. 207–11.
- 5 For example, Armstrong, *Perception and the Physical World*, 29–31; Cornman, *Perception, Common Sense and Science*, 254–8.
- 6 See the attack on power and force and other kinds of necessitation in Hume's *Treatise*, 155ff.
- 7 The best modern statement of the Humean position on *a priori* probabilities is found in Ayer's *Probability and Evidence*, 27–53.
- 8 Armstrong and McCall, 'God's lottery'. As this is a natural part of the argument from design, it is particularly good to find David Armstrong on the side of the angels.
- 9 Popper, *Realism and the Aims of Science*, 83.
- 10 Harré, *The Principles of Scientific Thinking*, 305.
- 11 Mackie, *Problems from Locke*, 25.
- 12 Ibid.
- 13 *Naming and Necessity*, 132ff.
- 14 This was Ayer's position in *Language, Truth and Logic*. It is, in general, a positivist theory.
- 15 Although he abandoned his positivism in many ways, Ayer retained this attitude towards ontology. It is explicit in the chapter 'The construction of the physical world' in *Central Questions of Philosophy* and in his 'Reply to John Foster' in the 'Library of Living Philosophers' volume dedicated to Ayer.
- 16 See Chapters IV and V above.
- 17 *Perception and the Physical World*, 65–7.
- 18 Hume, *Treatise*, I, iv, II, 'Of scepticism with regard to the senses'.
- 19 *Treatise*, 211.
- 20 Chapter VI, pp. 152–9.

Bibliography

- Annas, J., *Hellenistic Philosophy of Mind*, Berkeley, University of California Press, 1992.
- Annas, J. and Barnes, J., *The Modes of Scepticism*, Cambridge, Cambridge University Press, 1985.
- Anscombe, G.E.M., 'The intentionality of sensation; a grammatical feature', in R.J.Butler (ed.) *Analytical Philosophy*, second series, Oxford, Basil Blackwell, 1965.
- Aristotle, *De Anima* and *De Sensu* in vol. 1 of *The Complete Works of Aristotle: Revised Oxford translation*, (2 vols), ed. J.Barnes, Princeton, Princeton University Press, 1984.
- Armstrong, D.M., *Perception and the Physical World*, London, Routledge & Kegan Paul, 1961.
- Armstrong, D.M., *A Materialist Theory of the Mind*, London, Routledge & Kegan Paul, 1968.
- Armstrong, D.M., *Universals and Scientific Realism*, (2 vols), Cambridge, Cambridge University Press, 1978.
- Armstrong, D.M. and McCall, S., 'God's lottery', *Analysis*, 49, 223–4, 1989.
- Austin, J.L., *Sense and Sensibilia*, Oxford, Clarendon Press, 1962.
- Ayer, A.J., *Foundations of Empirical Knowledge*, London, Macmillan, 1940.
- Ayer, A.J., *Language, Truth and Logic*, (2nd edn), London, Gollancz, 1946.
- Ayer, A.J., 'Can there be a private language?', in *Concept of a Person and Other Essays*, London, Macmillan, 36–51, 1963.
- Ayer, A.J., *The Origins of Pragmatism*, London, Macmillan, 1968.
- Ayer, A.J., *Probability and Evidence*, London, Macmillan, 1972.
- Ayer, A.J., *Central Questions of Philosophy*, London, Wiedenfeld, 1973.
- Ayers, M. *Locke*, (2 vols), London, Routledge, 1991.
- Barnes, J., *The Toils of Scepticism*, Cambridge, Cambridge University Press, 1990.
- Barnes, W.F.H., 'The myth of sense-data', *Proceedings of the Aristotelian Society*, 45, 89–117, 1944–5.
- Bell, D., *Husserl*, London, Routledge, 1990.
- Berkeley, G., *Philosophical Works*, ed. M.R.Ayers, London, Dent, 1975.

Bibliography

- Bradley, F.H., *Appearance and Reality*, Oxford, Clarendon Press, 1930.
- Borst, C., *The Mind-Brain Identity Theory*, London, Macmillan, 1970.
- Brentano, F., *Psychology from an Empirical Standpoint*, London, Routledge & Kegan Paul, 1973.
- Brentano, F., *Sensory and Noetic Consciousness*, ed. O.Kraus, London, Routledge & Kegan Paul, 1981.
- Broad, C.D., *Scientific Thought*, London, Routledge & Kegan Paul, 1923.
- Broad, C.D., 'Some elementary reflexions on sense-perception', *Philosophy*, 27, 3–17, 1952; also in Swartz, *Perceiving, Sensing and Knowing*, 29–48.
- Broad, C.D., 'The theory of sensa', 1923, in Swartz, *Perceiving, Sensing and Knowing*, 85–129.
- Burge, T., 'Individualism and the mental', in *Studies in Metaphysics*, P. Frence, T.Vehling, H.Wettstein (eds), Minneapolis, University of Minnesota Press, 1979.
- Burge, T., 'Individualism and psychology', *Philosophical Review*, 95, 3–45 1986.
- Chisholm, R., *Perceiving: A Philosophical Study*, Ithaca, NY, Cornell University Press, 1957.
- Churchland, P.M., *A Neurocomputational Perspective*, Cambridge, Mass., MIT Press, 1989.
- Cohen, S.M., 'Saint Thomas Aquinas on the immaterial perception of sensible forms', *Philosophical Review*, XCI, 193–209, 1982.
- Cornman, J.W., *Perception, Common Sense and Science*, New Haven, Yale University Press, 1975.
- Crane, T. (ed.) *The Contents of Experience*, Cambridge, Cambridge University Press, 1992.
- Dancy, J. (ed.) *Perceptual Knowledge*, Oxford, Oxford University Press, 1988.
- Danto, A., Morgenbesser, S. (eds) *Philosophy of Science*, New York, Meridian Books, 1960.
- Davidson, D., *Essays on Actions and Events*, Oxford, Clarendon Press, 1980.
- Davidson, D., *Inquiries into Truth and Interpretation*, Oxford, Clarendon Press, 1984.
- Davies, M., 'Individualism and perceptual content', *Mind*, 100, 461–84, 1991.
- Dennett, D., *Brainstorms*, Hassocks, Harvester, 1978.
- Dennett, D., *The Intentional Stance*, Cambridge, Mass., MIT Press, 1987.
- Dennett, D., 'True believers', in *The Intentional Stance*, Cambridge, Mass., MIT Press, 1987.
- Dennett, D., *Consciousness Explained*, London, Allen Lane, 1991.
- Descartes, R., *The Philosophical Writings of Descartes*, trans. J.Cottingham, R.Stoothof, D.Murdoch, 2 vols, Cambridge, Cambridge University Press, 1984–5.
- Dretske, F., *Seeing and Knowing*, London, Routledge & Kegan Paul, 1969.
- Dretske, F., *Knowledge and the Flow of Information*, Oxford, Basil Blackwell, 1981.

Bibliography

- Dreyfus, H.L., 'Husserl's perceptual *noema*', in Dreyfus and Hall, *Husserl, Intentionality and Cognitive Science*, 97–123.
- Dreyfus, H.L. and Hall, H. (eds) *Husserl, Intentionality and Cognitive Science*, Cambridge, Mass., MIT Press, 1984.
- Ducasse, C.J., 'Moore's refutation of idealism', in Schilpp, *The Philosophy of G.E.Moore*, 223–51.
- Dummett, M., 'Common sense and physics', in Macdonald, *Perception and Identity*, 1–40.
- Eddington, A.S., *The Nature of the Physical World*, Cambridge, Cambridge University Press, 1929.
- Feigl, H. 'The "mental" and the "physical"', in *Concepts, Theories and the Mind-Body Problem*, Minnesota Studies in the Philosophy of Science, II, H.Feigl, M.Scriven, and G.Maxwell (eds), Minneapolis, University of Minnesota Press, 1958.
- Firth, R., 'Sense-data and the percept theory': Part I, *Mind*, 58, 436–65, 1949; Part II, *Mind*, 59, 35–56, 1950.
- Føllesdal, D., 'Brentano and Husserl on intentional objects and perception', in *Husserl, Intentionality and Cognitive Science*, ed. Dreyfus, 31–41.
- Føllesdal, D., 'Husserl's notion of *noema*', in Dreyfus and Hall, *Husserl, Intentionality and Cognitive Science*, 73–80.
- Føllesdal, D., 'Husserl's theory of perception', in Dreyfus and Hall, *Husserl, Intentionality and Cognitive Science*, 93–6.
- Foster, J., *Ayer*, London, Routledge, 1985.
- Foster, J. and Robinson, H., *Essays on Berkeley*, Oxford, Clarendon Press, 1985.
- Ginnane, W.J., 'Thoughts', *Mind*, 69, 372–90, 1960.
- Goodman, N., *Fact, Fiction and Forecast*, (3rd edn), Indianapolis, Bobbs-Merrill, 1973.
- Goodman, N., *Ways of Worldmaking*, Indianapolis, Hackett, 1978.
- Grice, H.P., 'The causal theory of perception', *Proceedings of the Aristotelian Society*, Supp. vol. 35, 121–52, 1961.
- Hacker, P.M.S., *Appearance and Reality*, Oxford, Basil Blackwell, 1989.
- Hahn, L., (ed.) *The Philosophy of A.J.Ayer*, Illinois, Open Court, 1982.
- Haldane, J.J., 'Aquinas on sense perception', *Philosophical Review*, 92, 233–9, 1983.
- Hamlyn, D.W., *Sensation and Perception*, Routledge & Kegan Paul, 1961.
- Hardin, C.L., *Color for Philosophers*, Indianapolis, Hackett, 1988.
- Hardin, C.L., 'Colour and illusion', in Lycan, *Mind and Cognition*, 555–67.
- Harré, R., *The Principles of Scientific Thinking*, London, Macmillan, 1970.
- Hinton, J., *Experiences*, Oxford, Clarendon Press, 1973.
- Hinton, J., 'Visual experiences', *Mind*, 76, 217–27, 1967.
- Hirst, R.J., *The Problems of Perception*, London, George Allen & Unwin, 1959.
- Honderich, T. (ed.) *Morality and Objectivity; a Tribute to J.L.Mackie*, London, Routledge & Kegan Paul, 1985.

Bibliography

- Hopkins, J., 'Wittgenstein and Physicalism', *Proceedings of the Aristotelian Society*, 75, 121–6, 1974–5.
- Hume, D., *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, eds L.A.Selby-Bigge and P.H.Nidditch, Oxford, Clarendon Press, 1975.
- Hume, D., *A Treatise of Human Nature*, ed. P.H.Nidditch, Oxford, Clarendon Press, 1978.
- Husserl, E., *Logical Investigations*, trans. J.Findlay, New York, Humanities Press, 1970.
- Husserl, E., *Ideas*, trans. W.R.Boyce Gibson, New York, Collier Books, 1972.
- Jackson, F., 'On the adverbial analysis of visual experience', *Metaphilosophy*, vol. 6, 127–35, 1975.
- Jackson, F., *Perception*, Cambridge, Cambridge University Press, 1977.
- Jackson, F., 'Epiphenomenal qualia', *Philosophical Quarterly*, 32, 127–36, 1982.
- Jackson, F., 'What Mary didn't know', *Journal of Philosophy*, 83, 291–5, 1986.
- Kenny, A., (ed.) *Aquinas*, London, Macmillan, 1969.
- Kirk, G.S., Raven, J.E. and Schofield, M., *The Presocratic Philosophers*: 2nd edn, Cambridge, Cambridge University Press, 1983.
- Kripke, S., *Wittgenstein and Rule-Following*, Oxford, Clarendon Press, 1976.
- Kripke, S., *Naming and Necessity*, Oxford, Blackwell, 1980.
- Lloyd, A.C., *Form and Universal in Aristotle*, Liverpool, Francis Cairns, 1981.
- Locke, D., *Perception and our Knowledge of the External World*, New York, George Allen & Unwin, 1967.
- Locke, J., *An Essay Concerning Human Understanding*, ed. P.H.Nidditch, Oxford, Clarendon Press, 1975.
- Lockwood, M., *Mind, Brain and the Quantum: The Compound 'I'*, Oxford, Basil Blackwell, 1989.
- Lycan, W., *Consciousness*, Cambridge, Mass., MIT Press, 1987.
- Lycan, W. (ed.) *Mind and Cognition: A Reader*, Oxford, Basil Blackwell, 1990.
- Macdonald, C., 'Weak externalism and mind-body identity' *Mind*, 99, 387–404, 1990.
- Macdonald, G.F. (ed.) *Perception and Identity: Essays Presented to A.J. Ayer*, London, Macmillan, 1979.
- McDowell, J., 'Criteria, defeasibility and knowledge', *Proceedings of the British Academy*, 1982, 455–79, London, Oxford University Press, 1983.
- McDowell, J., 'Values and secondary qualities', in Honderich, *Morality and Objectivity*, 110–29.
- Mace, C.A. (ed.) *British Philosophy in Mid-Century*, London, George Allen & Unwin, 1957.
- McGinn, C., *Mental Content*, Oxford, Basil Blackwell, 1989.
- Mackie, J.L., *Problems from Locke*, Oxford, Clarendon Press, 1976.
- McTaggart, J.M.E., *The Nature of Existence*, Cambridge, Cambridge University Press, 2 vols, 1921 and 1927.
- Malebranche, N., *Philosophical Selections*, ed. S.Nadler, Indianapolis, Hackett.

Bibliography

- Moore, G.E., 'The refutation of idealism', *Mind*, 12, 433–53, 1903.
- Moore, G.E., 'Some judgements of perception', *Proceedings of the Aristotelian Society*, 19, 1–29, 1918–19.
- Moore, G.E., 'A reply to my critics', in Schilpp, *The Philosophy of G.E. Moore*, 533–677.
- Moore, G.E., 'Visual sense-data', in Mace, *British Philosophy in MidCentury*, 203–11.
- Peacocke, C., *Sense and Content*, Oxford, Clarendon Press, 1983.
- Penfield, W., *The Excitable Cortex in Conscious Man*, Liverpool, Liverpool University Press, 1958.
- Pitcher, G., *A Theory of Perception*, Princeton, Princeton University Press, 1970.
- Place, U.T., 'Is consciousness a brain process?', *British Journal of Psychology*, 47, 44–50, 1956.
- Popper, K., *Realism and the Aims of Science*, ed. W.W.Bartley, III London, Hutchinson, 1983.
- Price, H.H., *Perception*, London, Methuen, 1932.
- Price, H.H., 'Appearing and appearances', *American Philosophical Quarterly*, 1, 3–19, 1964.
- Pritchard, H.A., *Knowledge and Perception*, Oxford, Clarendon Press, 1950.
- Putnam, H., 'The meaning of "meaning"', in *Mind, Language and Reality*, Cambridge, Cambridge University Press, 215–71, 1975.
- Reid, T., *The Works of Thomas Reid*, ed. Sir William Hamilton, Edinburgh, Maclachlan & Stewart, 1858.
- Robinson, H.M., 'The irrelevance of intentionality to perception', *Philosophical Quarterly*, 24, 300–15, 1974.
- Robinson, H.M., *Matter and Sense*, Cambridge, Cambridge University Press, 1982.
- Robinson, H.M., 'The general form of the argument for Berkeleyian idealism', in Foster and Robinson, *Essays on Berkeley*, 163–86.
- Robinson, H.M., 'A dualist account of embodiment', in Smythies and Beloff, *The Case for Dualism*, 43–57.
- Robinson, H.M., 'The objects of perceptual experience', *Proceedings of the Aristotelian Society*: supp. vol. 64, 151–66, 1990.
- Robinson, H.M., 'The flight from mind', in Tallis and Robinson, *The Pursuit of Mind*, 9–25.
- Robinson, H.M., 'Experience and externalism: a reply to Peter Smith', *Proceedings of the Aristotelian Society*, 92, 221–3 1992.
- Robinson, H.M., 'Physicalism, externalism and representation', in Wright, *New Representationalisms*, 103–14.
- Robinson, H.M., 'The anti-materialist strategy and the knowledge argument', in *Objections to Physicalism*, ed. H.M.Robinson, 159–83.
- Robinson, H.M. (ed.) *Objections to Physicalism*, Oxford, Clarendon Press, 1993.
- Russell, B., *Analysis of Matter*, London, Kegan Paul, 1927.

Bibliography

- Ryle, G., *The Concept of Mind*, London, Hutchinson, 1949.
- Schilpp, P.A. (ed.) *The Philosophy of G.E.Moore*, Chicago, Northwestern University Press, 1942.
- Searle, J., 'Minds, brains and programs', *Behavioral and Brain Sciences*, 3, 417–58, 1980.
- Searle, J., *Intentionality*, Cambridge, Cambridge University Press, 1983.
- Searle, J., 'Is the brain a digital computer?', *American Philosophical Association Proceedings*, 64 (3), 21–37, 1991.
- Sheehan, P., 'Aquinas on intentionality', in *Aquinas*, ed. A.Kenny, 307–21.
- Shoemaker, S., 'Functionalism and qualia', *Philosophical Studies*, vol. 27, 291–315, 1975; also in *Identity, Cause and Mind*, 184–205.
- Shoemaker, S., *Identity, Cause and Mind*, Cambridge, Cambridge University Press, 1984.
- Shoemaker, S., 'Absent qualia are impossible: a reply to Ned Block', *Philosophical Review*, vol. 90, 581–99, 1981; also in *Identity, Cause and Mind*, 309–406.
- Smart, J.J.C., 'Sensations and brain processes', *Philosophical Review* 68, 141–56, 1959.
- Smart, J.J.C., *Philosophy and Scientific Realism*, London, Routledge & Kegan Paul, 1963.
- Smith, A.D., 'Berkeley's central argument against substance', in Foster and Robinson, *Essays on Berkeley*, 37–57.
- Smith, D.W., 'Husserl on demonstrative reference and perception', in Dreyfus and Hall, *Husserl, Intentionality and Cognitive Science*, 193–213.
- Smith, P., 'On the objects of perceptual experience', *Proceedings of the Aristotelian Society*, 91, 191–6, 1991.
- Smythies, J.R. and Beloff, J. (eds) *The Case for Dualism*, Charlottesville, University of Virginia Press, 1989.
- Snowdon, P., 'Perception, vision and causation', *Proceedings of the British Academy*, vol. 68, 455–79, 1982; also in *Perceptual Knowledge*, ed. J.Dancy, 192–208.
- Sorabji, R., 'From Aristotle to Brentano: the development of the concept of intentionality', in *Oxford Studies in Ancient Philosophy: Supp. vol.* 227–59, 1991.
- Sorabji, R., 'Intentionality and the physiological processes: Aristotle's theory of perception', in A.Rorty and M.Nussbaum (eds), *Aristotle's De Anima*, Oxford, Clarendon Press, 195–255, 1991.
- Stebbing, S., *Philosophy and the Physicists*, London, Methuen, 1937.
- Strawson, P., 'Perception and its objects', in *Perception and Identity*, ed. Macdonald, 41–60.
- Swartz, R.J. (ed.) *Perceiving, Sensing and Knowing: A book of readings from twentieth century sources in the philosophy of perception*. Garden City, New York, Doubleday, 1965.
- Tallis, R. and Robinson, H. (eds) *The Pursuit of Mind*, Manchester, Carcanet Press, 1991.

Bibliography

- Theophrastus, *De Sensu*, in *Theophrastus and the Greek Physiological Psychology*, by G.M.Stratton, London, George Allen & Unwin, 1917.
- Tye, M., 'The adverbial approach to visual experience', *Philosophical Review*, 93, 195–225, 1984.
- Tye, M., *The Metaphysics of Mind*, Cambridge, Cambridge University Press, 1989.
- Walker, R.C.S., *The Coherence Theory of Truth*, London, Routledge, 1989.
- Warner, R., 'Incorrigibility' in *Objections to Physicalism*, ed. H.M. Robinson, 182–205.
- Wittgenstein, L., *Philosophical Investigations*, Oxford, Basil Blackwell, 1954.
- Wood, M., 'Universal and particular forms: Aristotle's *Metaphysics*', in *Oxford Studies in Ancient Philosophy*: supp. vol., 41–56, 1991.
- Wright, E., (ed.) *New Representationalisms*, Avebury, Avebury Press, 1993.
- Yolton, J.W., *Perceptual Acquaintance from Descartes to Reid*, Oxford, Basil Blackwell, 1984.

Index

- a priori* probabilities 216–19
- adverbialism 136, 174–83
- Anscombe, G.E.M. 167–72
- appearance: judgemental theory of
 - 50–4; multiple location theory of 44–8; relational theory of 48–50
- Aquinas, St Thomas 6–10
- Armstrong, D.M. 63, 94, 121, 127, 144–5, 190, 230
- Aristotle 4–10, 19
- associationism 19
- atomism 15, 37, 59, 84
- Austin, J.L. 32, 51–3, 55
- Ayer, A.J. 47, 102–5, 202

- Barnes, W.H.F. 42–5
- Berkeley, George 1, 3, 12, 13, 14, 18, 19, 20, 33, 35, 36, 88, 206, 214, 226, 229–31
- blind spots 202–5
- Boyle, R. 15
- Brentano, F. 7, 8, 21–4, 25, 26, 29, 45, 178
- Bradley, F.H. 41
- Broad, C.D. 37, 41, 51, 54, 57
- Burge, T. 137, 145

- Charles, King and Martyr 81, 87
- Churchland, Patricia 125–7
- Churchland, Paul 125–7
- cognitivism 122–5
- Cohen, S. 7

- Davidson, D. 158

- Dennett, D. 111, 134, 158, 193–202
- depth 205–7
- Descartes, R. 11, 14, 15, 16, 17, 18, 19, 33, 34, 35, 36, 88, 206, 214, 219, 226
- disjunctive analysis 152–9, 163, 173, 237–8
- dispositional direct realism 71–3, 79
- dreaming 35
- Dretske, F. 121, 127
- Ducasse, C.J. 182–3
- Dummett, M. 77–9

- Eddington, A.S. 75
- Evans, G. 27
- evil demon 17, 35
- externalism 136–50, 173, 179, 181

- Feigl, H. 122, 185
- Firth, R. 28, 29
- Frege, G. 25, 27
- functionalism 128–36

- Galileo 59
- Ginnane, W. 22
- Goodman, N. 113, 196
- Grice, H.P. 52–4

- Hacker, P.M.S. 54–6, 62–3, 74, 103
- hallucination 17, 18, 142, 144, 153; argument from 87–8, 153
- Hardin, C.L. 64
- Harré, R. 220

Index

- Henry VIII 165
 Hinton, M. 154, 156–8, 160
 Hopkins, J. 98
 Hume, D. 1, 3, 12, 13, 14, 19, 20, 21, 33, 34, 35, 36, 215, 216, 220, 226, 231–6
 Husserl, E. 24–7, 29, 45
- idealism 3, 213
 illusion: argument from 31–58, 74, 236; definition of 54
 imagism 14
 impenetrability 219
 information 63
 intentionality 2, 10, 18, 20, 21–7, 32, 72, 89, 158, 163–86;
 intentional objects 2, 8, 25, 43, 55, 71, 79, 95, 131, 136, 179, 237
- Jackson, F. 180, 199, 202
- Kant, I. 21, 71, 229–30
 Koler's *phi* phenomenon 196–7
 Kripke, S. 27, 95, 98, 105–16, 149, 221
- Leibnitz's Law 32, 81
 Locke, D. 153, 154
 Locke, J. 1, 3, 5, 11–14, 16, 17, 18, 19, 20, 33, 35, 60, 70, 75, 85, 214, 219
 Lucretius 59
 Lycan, W. 178
- Macdonald, C. 145–7
 McGinn, C. 145
 Mackie, J. 220
 McTaggart, J.M.E. 41
 Marr, D. 137
 Martin, M. 140–1, 143, 144, 145
 Meinong, A. 25, 178
 memory 117–18
 Mill, J.S. 3, 27
 Molyneux's problem 207–11, 214
 Moore, G.E. 37–41, 54, 182, 183, 184
- Muller-Lyre lines 31
- noema* 25–7
 nominalism 36, 208
- percept theory 27–30, 187, 188
 perspectivalism 77–9
 phantasm 4, 6, 11, 12, 15, 17, 19, 208
 Phenomenal Principle 31–2, 36, 49, 54–6, 58, 186
 phenomenalism 3, 176, 213, 222–38
 Pitcher, G. 121, 154–6
 Place, U.T. 119
 Popper, K. 219
 presentational nature of experience 23, 123
 Price, H.H. 29, 40, 41, 45–9, 54
 primary qualities 60, 61, 63, 85, 207
 principle of minimal empiricism 122–4, 128–9
 privacy 71, 73, 91–118
 projectivism 71
- recognitional conception 97–105
 representative realism 3, 6, 9, 176, 213–22
 Reid, T. 19–21, 45
 Russell, B. 27, 214
 Ryle, G. 55, 108
- science: argument from 74–80
 Searle, J. 150
 secondary qualities 3, 15, 16, 21, 22, 29, 37, 49, 55, 56, 85, 95, 207, 214; argument from 59–74
 sense-data: definition of 1–2; maximal position concerning 188, 202–7
 Shoemaker, S. 201
 Smart, J.J.C. 94, 127
 Smith, A.D. 18
 Smith, P. 139
 solidity 76–7, 219, 220
 Strawson, P.F. 27, 77–9

Index

- topic neutrality 63, 94, 109, 137,
139, 173, 190, 221, 243
- Tye, M. 178, 181
- universals 13; resemblance theory
of 63
- Walker, R.C.S. 115–16
- Wiggins, D.R.P. 225
- Wittgenstein, L. 55, 74, 91–118,
120, 229
- Yolton, J. 18, 19